

Top 10 PowerShell ISE Tips

Windows IT Pro

A PENTON PUBLICATION

OCTOBER 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Deploy DirectAccess in Windows Server 2012


PowerShell's Select-Object

Implement Continuously
Available File Shares in
Windows Server 2012

System Center 2012 SP1
Virtual Machine Manager

Sean Deuby on
Federated Logon





Windows IT Pro Store

eLearning Classes

eBooks

On-Demand Training

In-Person Training

Posters

Videos

Plus you can **RENEW** your subscription or
UPGRADE to VIP membership while
you're there!

Stop by the store today!

WindowsITPro

FREE Newsletters!

**Not your average
Newsletters!**

WinInfo Daily UPDATE

Paul Thurrott covers the entire Windows universe with reviews, commentary, analysis, and tips. Delivered daily.

Windows IT Pro UPDATE

Windows industry news, products, FAQs, tips, and resources for IT professionals. Delivered weekly.

Cloud & Virtualization UPDATE

Get the latest news, blogs and analysis to help you determine your organization's cloud and virtualization strategy. Delivered weekly.

Exchange and Outlook UPDATE

News, strategies, products, and developments in Exchange Server and Outlook messaging. Delivered weekly.

Security UPDATE

Learn about Windows security risks, attacks, and how to fix or avoid them. Includes security alerts! Delivered bi-weekly.

Dev Pro UPDATE

Topics for Microsoft platform developers: ASP.NET, .NET Framework, Silverlight, mobile, and SQL Server development. Delivered weekly.

SQL Server Pro UPDATE

The latest news, products, and developments for SQL Server DBAs and developers. Delivered weekly.

SharePoint Pro UPDATE

SharePoint for IT professionals and developers – weekly tips, news, and how-to's. Delivered weekly.

**subscribe today at
windowsitpro.com/manage-newsletters**

WindowsITPro

COVER STORY ▼

Deploying DirectAccess 21 in Windows Server 2012

— John Savill

DirectAccess offers a great end-user experience and powerful management capabilities—far beyond what's possible with a manually initiated VPN connection.

Features

- 33 PowerShell Basics: Select-Object**
Jeffery Hicks
- 40 Microsoft System Center 2012 SP1 Virtual Machine Manager User Roles**
Damir Dizdarevic
- 50 Windows Server 2012: Implement Continuously Available File Shares**
Michael Otey

Products

- 75 New & Improved**

Interact

- 70 Ask the Experts**

In Every Issue

- 79 Advertiser Directory**
- 79 Directory of Services**
- 79 Vendor Directory**

Chat with Us



Facebook



Twitter



LinkedIn

Columns



6

Need to Know

Ballmer Exits, Windows 8.1 Is Finalized

Paul Thurrott



12

Windows Power Tools

Installing Printers with PowerShell

Mark Minasi



15

Top 10

Top 10 Tips for Using PowerShell ISE

Michael Otey



19

Enterprise Identity

Does Federated Logon Pass Consumer Testing?

Sean Deuby

Editorial

Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
IT Community Manager: Rod Trent
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialist: Karly Prickett

Advertising Sales

Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service • 800-793-5697
Vice President, User Marketing &
Marketing Analytics: Tricia Syed
Marketing Director: Amy Connell

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, October 2013, Issue No. 230,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton. Copyright ©2013 Penton. All rights reserved. No
part of this publication may be reproduced or distributed
in any way without the written consent of Penton.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2013, Penton, all rights reserved.
These programs may not be reproduced or distributed
in any form without permission in writing from the
publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton, under license from owner.
Windows IT Pro is an independent publication not
affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

Ballmer Exits, Windows 8.1 Is Finalized



**Paul
Thurrott**

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



When Microsoft CEO Steve Ballmer unexpectedly announced in August that he would leave Microsoft within 12 months, he left more questions than answers. Ballmer's resignation should have a profound effect on everyone who uses Microsoft's products, and it will be interesting to see how his departure affects the current product pipeline—including the recently finalized Windows 8.1.

Ballmer Is Out

After more than 13 years as Microsoft's CEO, Steve Ballmer announced in late August that he would leave the company within 12 months, after helping the firm's board of directors find a suitable replacement. [Ballmer's goodbye letter to employees](#) explained the timing—he said Microsoft needed a leader who would be there long enough to see the company through its transition to a devices and services firm, and he had originally intended to leave earlier, in 2018. But his letter was short on details about why he was leaving his position as CEO. Put simply, Ballmer's tenure as CEO is decidedly mixed.

On the one hand, Ballmer presided over a period of dramatic economic expansion for Microsoft. As he noted in his letter to employees, the firm grew from annual revenues of \$7.5 million in fiscal 2000 to nearly \$78 billion in fiscal 2013. He was the 30th employee in a company that now employs almost 100,000 people. The firm claims more than 1 billion users worldwide, and Ballmer says that Microsoft has “delivered more profit and cash return to shareholders than virtually any other company in history.” There's nothing small about Microsoft.

On the other hand, Microsoft has been dogged by a series of strategic missteps that in many ways define Ballmer's era at the firm. Windows Millennium Edition, the aborted “Longhorn” project, Windows

Vista, and now Windows 8/Windows RT all happened on his watch. Dogged in part by antitrust oversight, Microsoft missed key market shifts and became a follower rather than a leader in important markets such as digital music, smartphones, tablets, and cloud computing. It was always a day late and a dollar short.

What protected Microsoft from imploding and, I think, hid the ills for too long, was that the firm's dominant product lines—Windows, Office, and Windows Server—kept generating record profits and revenues long after its fastest-moving competitors had moved on to up-and-coming markets that Microsoft should have seamlessly side-stepped into as well. (You can read more about this theory in [“Assessing the Ballmer Years.”](#))

Microsoft tried to move, under Ballmer, but always under the auspices of protecting Windows (and Office and Server) at all costs. We saw me-too products such as Zune, Windows Mobile 6.5, and Windows Live Search/Bing. When users embraced Gmail, Microsoft retrofitted Hotmail as Outlook.com. When users embraced the web, Ballmer tried to buy Yahoo!, his one truly outrageous mistake. (And a barely avoided disaster.) When users embraced cloud storage, Microsoft gave us SkyDrive. When the Apple iPad took over, we got Surface. Again and again, Microsoft let others blaze trails, then belatedly followed them after a market proved valuable.

But the product that best parallels the problem of the Ballmer years is Windows 8 (which includes the pointless Windows RT). Faced with an exodus of users, developers, and mindshare to mobile computing platforms such as iOS (iPhone/iPad) and Android, Microsoft had two courses it could take with Windows. It could simply continue to develop future iterations of the classic desktop OS while creating a purely mobile platform on the side, much as Apple did with Mac OS X and iOS, respectively. Or it could do what it's always done: Protect Windows at all costs and, in this case, simply build mobile platform features into Windows. Microsoft, after all, exults in the malleability of Windows.

Of course, it chose the latter path. Critics will point to this decision as a mistake and proof that Ballmer's long-term strategy was a mistake. But here's an inconvenient truth: Had Microsoft created a "Metro OS" or whatever, separately, for mobile devices, that system would have sunk in the market just as badly as has Windows 8, if not worse. As with Windows Phone before it, there just isn't much demand for yet another mobile platform, not when both Android and iOS have hundreds of thousands of apps and established ecosystems. Desktop Windows, meanwhile, would have continued its inevitable decline, racing to become the smallest of the three major mainstream computing markets.

But in melding the new Metro platform onto Windows, Microsoft has, in effect, forced all Windows customers to deal with this new mobile OS whether they want it or not. This has created an unprecedented backlash, triggering the development of a refined version of the OS, Windows 8.1.

Windows 8.1 softens the transition between the desktop and Metro and makes it possible for users to stick to the environment they prefer. Often described as a combination service pack/feature pack, Windows 8.1 is better seen as an apology, a mulligan aimed at easing friction in the user base. And as with the backlash that accompanied its release—it even sank the beautiful Surface hardware—this kind of retreat is itself unprecedented.

Give Ballmer some credit: Though the current quarter could indeed be abysmal by Microsoft standards, Ballmer never ran the company into the ground. Microsoft has the financial resources, if not the time, to make yet another comeback. The question is whether Ballmer's successor will continue down the company's current path—it describes itself now as a maker of "devices and services," though it has precious few success stories in either category—or tread a new path.

The issue here is the Microsoft board of directors. Led by company cofounder Bill Gates, the board is unlikely to take the harsh and necessary steps of really remaking Microsoft. And in its public

statements at the time of Ballmer's exit announcement, the board reiterated its support for the company's current strategy. So it's very likely that Microsoft will simply hire from within—or, intriguingly, bring back a previous executive such as Paul Maritz or Stephen Elop to right the ship.

But it's pretty clear that a company of Microsoft's size can coast for a long, long time without actually getting the strategy right. Perhaps the time to make real change is before such change is forced on the company. On that note, I'd like to see Microsoft hire an external candidate, as Ford did when it hired Boeing's Alan Mulally to rescue that company. Microsoft needs someone objective to rate its current path and determine whether further and drastic changes are required. I suspect that they are.

Windows 8.1 Finalized

Microsoft finalized Windows 8.1 on August 23, 2013, and will release it to customers online and via a new generation of PCs and devices (as of this writing) on October 17, 2013. This release is a bit of a cipher. It's an interim update for both Windows 8 and Windows RT, a sort of combination feature pack and service pack, as I previously noted. But it's also, in effect, a new version of Windows, and Microsoft uses the terms Windows 8.1 and Windows RT 8.1 to differentiate this new version from the initial versions delivered last year.

That Microsoft would like to distance itself from Windows 8 is completely understandable. The slowest-selling version of Windows in modern times, Windows 8 has undone all of the good will that Microsoft engendered with Windows 7, and then some. That it arrived at a time during which consumers were embracing simpler alternative mobile platforms is, of course, not coincidental.

But while Microsoft likes to brag that Windows 8.1 shows what the Windows team can accomplish in just one year, I think the lesson here is quite different: The firm should have simply waited until this release to ship anything. Windows 8.1 is a much more complete and

mature product than its predecessor, and it's much more respectful to the billion-plus users out there who use Windows with traditional, non-touch hardware.

So we see the much-ballyhooed return of the Start button in this release, which should smooth some rustled feathers, though I never saw its absence as an issue. No, Microsoft won't let you go back to the old Start menu, but it has made other concessions to typical PC users that should be appreciated.

You can, for example, boot right to the desktop, skipping the full-screen Start screen. And you can configure the system to display a desktop-oriented version of the All Apps screen instead of the Start screen when you use any of the usual methods to invoke that interface. All Apps works a bit more like the old Start menu.

There are deeply hidden controls to remove many other Metro interfaces, including the silly Switcher app-switching bar, and a partial remedy for the much-loathed Charms. The point here is that desktop users should be able to stick with the Windows desktop most of the time, a huge improvement over the initial version of Windows 8.

Conversely, Microsoft is also making it easier for Windows tablet users to use the touch-friendly Metro world. In this release, for example, most of the common system settings that were previously accessible only from the desktop-based Control Panel can now be found in the Metro-centric PC Settings.

Those who want to avoid the desktop entirely—and I'm told they exist—can mostly do so. (Office, of course, is currently available only in desktop form, as are the world's most popular Windows applications, Google Chrome and Apple iTunes.)

In this new Metro world, the capabilities of the built-in apps have gotten better. The Mail app now supports drag-and-drop for both touch and mouse, turning it into a usable solution. The bundled Bing apps—with updated versions of News, Travel, Sports, Finance, and Maps, plus new apps such as Food & Drink and Health & Fitness—are surprisingly good and show how well Metro-based content solutions can work.

With this release, Microsoft is also embracing a bundling strategy that might seem superficially similar to the bundling activities that got it in antitrust trouble 15 years ago. But Microsoft's lack of success in mobile has created a different environment, and the firm is right to push its other brands in Windows.

Thus, you'll see a lot of Bing in this release, not just in the aforementioned apps, but also in the stunning new Search experience. And Skype is bundled with the OS, just as Messenger was back in the day. So, too, are Xbox-branded Music, Video, and Game experiences, each upgraded. SkyDrive is there, too.

Furthermore, Microsoft is starting to bundle Office with more Windows PCs. With the initial wave of Windows 8 releases, consumers could pick up a free copy of Office Home & Student 2013 RT with each Windows RT device, and in Windows 8.1, Outlook RT has been added to the mix. Now anyone who purchases a Windows 8-based mini-tablet will also get a free copy of Office Home & Student 2013 (albeit the non-RT versions sans Outlook), too. After spending the past decade patiently explaining to people that Windows and Office aren't the same thing, I'm finding the lines are really starting to blur.

Ultimately, Windows 8.1 is exactly what you think it is: a better version of Windows 8. The question, however, is whether the evolutionary changes in this release warrant a reassessment of the platform. Will users embrace Windows 8.1 after ignoring its predecessor?

Honestly, I don't think so, and while a coming generation of hardware will certainly help, it's not clear whether Microsoft's vision of the mobile computing future has yet aligned with what users expect of Windows. And that's a long-term issue that Steve Ballmer's successor will need to address. ■

Installing Printers with PowerShell



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email



Twitter



Website



In “Managing Printers Gets Easier in Windows 8,” I observed that many folks have decided to skip Windows 8 thanks to its (ahem) less-than-sterling reputation—but those who *have* decided to adopt Windows 8 are discovering a bevy of hidden treasures. Among those gems are 20 useful printer-related cmdlets. In that article, I showed how the *add-printerdriver* cmdlet lets you designate one or more printer drivers as essentially “safe for the non-admin user to use,” paving the way for a cmdlet that I’ll cover this time: *add-printer*.

As you’ve probably already guessed, *add-printer* lets you install a printer. That doesn’t sound exactly scintillating—except for two things. First, *add-printer* is a command-line tool, which makes automating it easy. Second (and here’s the really nice part), non-administrative users can run it, which means that automating it is as simple as putting the cmdlet in a user’s logon script. (If that sounds troubling, remember that a user can’t create a printer unless you approve that printer’s driver.)

Add-Printer Options

The *add-printer* cmdlet has many options, but here’s the basic syntax:

```
add-printer -name <name> -drivername <driver name>
           -port <port name>
```

The first option is just the name you want to appear in Devices and Printers (e.g., “Upstairs laser”). The driver name is the same name you used in *add-printerdriver* (e.g., Brother HL-4040CDN Series, HP Deskjet 5700 Series, Dell Color Laser 1320c). But what about the port name? You’ve seen the PowerShell nouns *printer* and *printerdriver*, but there’s one more to learn: *printerport*.

You can see your system's printer ports by typing

```
get-printerport
```

You'll probably get a lot of results. Figure 1 shows a few from my system. I created the port named ToColor. I had a network-attached printer on my intranet—a Dell 1320c sitting on 10.50.50.50—that I wanted to connect to a workstation.

Name	ComputerName	Description	PortMonitor
----	-----	-----	-----
COM1:		Local Port	Local Monitor
FILE:		Local Port	Local Monitor
LPT1:		Local Port	Local Monitor
nul:		Local Port	Local Monitor
ToColor:		Standard TCP/IP Port	TCPMON.DLL
USB001		Virtual printer p...	Dynamic Print Mon...

So, I had this much of the *add-printer* cmdlet so far:

```
add-printer -name "ColorDell" -drivername "Dell
  Color Laser 1320c"
```

Figure 1
Printer Ports

What about *-port*? If the printer were directly USB-attached, I could use *-port USB001*, but I wouldn't because the system's built-in Plug and Play (PnP) infrastructure identifies the port name automatically. But if I want to connect to a network-attached printer, my system needs help. Assuming I know the IP address or DNS name of the printer, PowerShell's *add-printerport* lets me create a port to that address. So, if my Dell 1320C is at 10.50.50.50, I can create a network port to it with the command

```
add-printerport -name "ToColor" -printerhostaddress
  "10.50.50.50"
```

That's simple, because I only have to give the port a name ("ToColor") and an IP or DNS address. (You need only standard user privileges to

create printer ports with PowerShell.) With the printer port in place, any user with access to the Dell 1320C drivers can connect, like so:

```
add-printer -name "ColorDell" -drivename "Dell Color
  Laser 1320c" -port "ToColor"
```

What About Shared Printers?

The preceding two scenarios leave out the most common situation, which is connecting to a printer shared via a server. To connect to one of those, you need only the *-connectionname <UNC parameter>* option. Suppose I've connected that Dell 1320C printer ("ColorDell") to a server called Netdoor and shared that printer as \\netdoor\PL. My workstation could then connect with just *-connectionname*, as in

```
add-printer -connectionname \\netdoor\PL
```

In that case, I'd end up with a printer named \\netdoor\ColorDell and, yes, you read that right: PowerShell gives that connection a name that isn't blindly equal to the UNC path but rather blends the server's name and printer name perceived by the print server.

Printer Removal

What about disconnecting? If you're even a bit PowerShell-adept, you'll already know that PowerShell's verb for deleting or eliminating something is *remove*, and you'll have guessed that you delete a printer with *remove-printer*, as in

```
remove-printer "Downstairs Printer"
```

Always remember that although I often use uppercase and lowercase in my examples to make them a bit more readable, PowerShell is almost always case-insensitive. So, *remove-printer "DoWNstaiRS PRINTER"* would work just as well. ■

Top 10 Tips for Using PowerShell ISE

Microsoft's free ISE is essentially the standard tool for PowerShell development

If you're just getting started with [PowerShell](#), you'll probably be doing your work in the Integrated Scripting Environment (ISE). Although there are many third-party products that improve upon the features of ISE, Microsoft's free ISE is essentially the standard PowerShell development tool. Sure, you can edit your PowerShell scripts in just about any text editor, including the venerable Notepad, but ISE is a much more productive tool, providing you with the ability to use IntelliSense and color-coded syntax as well as edit, execute, and debug PowerShell scripts. In this column, I'll show you 10 tips to make your PowerShell development in ISE more productive.



Michael Otey

is senior technical director
for *Windows IT Pro* and
SQL Server Pro.



Email

① Put ISE on the Windows 8 Start Screen

Although PowerShell 3.0 and PowerShell ISE are both delivered with [Windows 8](#), there's no PowerShell ISE option on the Windows 8 Start screen or desktop, and if you search through Apps you won't find it. That doesn't mean that PowerShell ISE isn't there. It's hidden on the Administrative menu, which isn't displayed by default. To add the Administrative menu and the PowerShell ISE option to the Windows 8 Start screen, open the Windows 8 Settings charm, choose the Tiles option, then move the *Show administrative tools* slider to Yes.

② Set the Execution Policy

Oddly, although ISE is clearly oriented toward developing scripts, it does nothing to change PowerShell's default script execution policy,

which doesn't allow scripts to run. The default PowerShell execution policy is set to Restricted. To allow ISE to run PowerShell scripts, go to the Console pane and enter the following command:

```
Set-ExecutionPolicy RemoteSigned
```

③ Open Multiple Tabs

One of the things that makes ISE so much more powerful than Notepad is that it lets you open multiple tabs and work on multiple scripts at the same time. Unlike Notepad, it doesn't isolate you in a single window. You can approximate this functionality with multiple Notepad windows, but then you lose color coding, IntelliSense, and code snippets. To open multiple tabs, use the File, New option or the File, Open option, and ISE will open a new tab in the Scripting pane.

④ Use Snippets

If you're not a developer, you might not know what code snippets are all about. Code snippets are prebuilt code blocks that you can insert into the Scripting pane to give you a head start in writing the correct PowerShell code. For example, if you want to use an If-Else statement but you don't remember the exact syntax, you can simply position your cursor where you want the If-Else statement to start and then press Ctrl + J or select Start Snippets from ISE's Edit menu. Doing so will display a dialog box with all the available snippets. As you scroll through the dialog box, a tooltip displays the actual PowerShell code that will be inserted.

⑤ Use Code Regions

Another feature that ISE provides to help you navigate your code is regions. Regions are collapsible sections of code indicated by a minus sign and an outline marker on the left side of the Script pane. ISE automatically creates regions for block structures, such as If-Else, For-Next, For-Each, and While loops. You can also create your own regions by

marking the start of the region using the `#region` tag, optionally followed by a name. You mark the end of the region by using the `#endregion` tag. A closely related feature is PowerShell's automatic brace matching. If you select a brace or parenthesis, ISE will automatically highlight the matching brace or parenthesis.

⑥ Use F1 PowerShell Help

As you might expect, ISE provides a lot of help for people who are just getting started with PowerShell. The Command Add-In pane on the right side of the screen can help you see the valid parameters for the various PowerShell cmdlets. The built-in F1 Help goes further by displaying a graphical pop-up window displaying Help for a selected PowerShell cmdlet. You can take advantage of the pop-up F1 Help by simply moving your cursor over a cmdlet that you want to display Help for and pressing F1.

⑦ Run Code

Although it might not be as full featured as some of the third-party PowerShell development products, ISE is completely capable of running and debugging PowerShell code. To run just part of a script, highlight the text you want to run and click the Run Script icon or press F5. Doing so will run just the selected code. To run the entire script, click the Run Script icon or press F5 without making a specific code selection. You'll see the results of the PowerShell code displayed in the Console pane.

⑧ Set Breakpoints

For a serious script developer, one of the most important features in ISE is its integrated debugger. You can use breakpoints to stop the execution of a given PowerShell script on a specific line. Breakpoints can be set on lines or variables. To toggle a breakpoint on a line, right-click on the line where you want the code execution to stop, then select Toggle Breakpoint from the context menu. Alternatively,

you can click on the line and select Toggle Breakpoint from the Debug menu. You can also use the PowerShell *set-psbreakpoint* and *get-psbreakpoint* cmdlets to set and view breakpoints. You can't set breakpoints on comment lines.

⑨ Single Step with the Debugger

Just as important as setting breakpoints is the ability to track the execution of your code by single stepping through the code. Single stepping through your code can help uncover logic problems that the code might have. The easiest way to single step is to press the F10 key after the code has halted on a breakpoint. You can also select Step Over from the Debug menu. If you have looping structures or functions that you want to step through, you can use F11 or Step Into from the Debug menu. Shift + F11 or Step Out will quickly exit the loop or function.

⑩ Examine Variables

Although stepping through your code is a valuable tool for uncovering logic errors, the ability to display the contents of variables is just as important. To display the contents of a variable, simply hover the mouse over any occurrence of the variable in the Script pane. You can also go to the Console pane and type in the variable name and press Enter. Of course, the execution of the script needs to be paused when you display a variable. ■

Does Federated Logon Pass Consumer Testing?

Sometimes the answer isn't only technology

I obviously spend a lot of time reading, writing, and listening to bright people talk about the complicated problem of Internet identity. The other day, I unexpectedly had a pop quiz on my ability to explain what the Internet identity community is trying to accomplish. And I had a tough audience: my wife.

Now, my wife is pretty computer savvy for a classical musician. (She'd have to be after all this time around me.) But, as with most everyone else, keeping track of a plethora of website user IDs and passwords makes her crazy. When she was contemplating signing up for LinkedIn, I noticed that "Sign In With Facebook" was an available option instead of creating a local LinkedIn account. I jumped on it, of course. [Federated identity!](#) [Internet single sign-on \(SSO\)!](#) [Just-in-time provisioning!](#) Then the depth of my challenge sunk in.

The Facebook Challenge

I had to succinctly explain (remembering that, to the rest of the world, identity is just a speed bump in the way of the end goal—in this case, using LinkedIn) why she'd want to go ahead and let Facebook share her data with another website. After all, Facebook has become infamous for violating accepted privacy practices by default, then backtracking when there's an outcry. To add pressure, I knew that if I convinced her and it was a poor experience, it'd be a long time before I could convince her to try again.

After quickly throwing out a number of possible motives why she'd want to use her Facebook creds (see "Federated identity! Internet SSO!" above), I settled on two. First, she wouldn't need to maintain a



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.



Email



Twitter

separate LinkedIn account. Second, although Facebook would supply some identity information, her Facebook password wouldn't go anywhere, because LinkedIn trusts Facebook to provide legitimate IDs. Somewhat skeptically, she agreed to continue the experiment. I was a little concerned about Facebook being the only choice of identity provider. Other identity providers (Google, for example) typically supply the relying party (LinkedIn) with the bare minimum of information, but Facebook shovels a lot of identity data at the relying party.

Sure enough, when she signed in with her Facebook account, the authorization dialog box told her that LinkedIn wanted to access her public profile, friends list, email address, work history, education history, and current city. End of experiment! She was having none of this “kitchen sink” approach. She quite reasonably wanted *control* over what information Facebook was providing to LinkedIn. Because she wasn't offered any choice in the matter (and wasn't even offered an explanation of why LinkedIn wanted this information), she opted out of the whole thing.

Game Over, Man! Game Over!

So much for the nirvana of federated logon. Trust has become a precious commodity on the web, and Facebook in particular has beaten up the consumer's trust in its approach to privacy. But Facebook isn't alone, of course; if the entire federated logon experience (aka “logon with”) isn't completely trustworthy and clearly explained to the user, it will be a very hard sell to get consumers to trust it. What's needed are clear explanations of *why* certain types of information are required. Why is it safer to use an existing user ID and password than creating a separate one? Why does the relying party even need these pieces of information from the identity provider? In this case, the answer to cloud identity confusion isn't technical. A little explaining goes a long way. ■

Deploying DirectAccess in Windows Server 2012

Installing and configuring DirectAccess is now painless

In Windows Server 2008 R2, installing and configuring DirectAccess is painful. It requires a complicated setup process that involves meeting public key infrastructure (PKI) requirements, getting server certificates, setting up a network location server, making sure the targets support IPv6, and using the Forefront Unified Access Gateway (UAG). In *Windows Server 2012*, installing and configuring DirectAccess is simple if you're using Windows 8 clients. Before I describe the vastly improved setup process, I'll take a step back and tell you what DirectAccess is and how it can help your organization.

What Is DirectAccess?

In my first job as a VAX/VMS systems administrator, I typically got to the office at 8:00 A.M. and left at 5:00 P.M. Those were the hours I interacted with the work systems (including email), and I never accessed the systems outside the office. Today, the concept of 8-to-5 office hours has disappeared, and the line between work and personal life has blurred. IT administrators and end users alike need to be able to access company systems and data all the time, so they always need connectivity.



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).



There are two traditional approaches for managing access to corporate systems from outside the corporate network:

- Access over web-based protocols, such as HTTP Secure (HTTPS). For example, this type of access is used for accessing Microsoft Exchange mailboxes over ActiveSync, accessing SharePoint sites that are published in a secure fashion to the web, and even using remote desktop connections by means of the Remote Desktop Gateway, which encapsulates the RDP traffic in HTTPS.
- VPN tunnels through the Internet between a machine and the corporate network. With this approach, users manually initiate a connection to the corporate network.

Using HTTPS is a great solution when it's available. It has the advantage of typically "just working" and is available on many different types of devices. However, the HTTPS approach doesn't work for many types of services, such as line of business (LOB) applications. And sometimes organizations don't want to use it, even if it's a viable option. For these situations, the traditional VPN approach can be used. But this approach also has challenges:

- The users must manually initiate the connection, which can be complex.
- The users are connected to the corporate network infrastructure only when they're in a VPN session. As a result, if the users don't connect often, their computers can't be managed for activities such as patching, policy updates, and software updates.

To provide another option, Microsoft introduced DirectAccess in Server 2008 R2 and Windows 7 (Enterprise and Ultimate editions). DirectAccess enables an always-on connection from the client to the corporate network, without any user action. When users are accessing an Internet resource, their regular Internet connection is used. When users are accessing a corporate resource, the DirectAccess tunnel is used, giving them transparent corporate resource access from

anywhere. To determine whether the target is a corporate resource, DirectAccess compares the target's DNS suffix against a Name Resolution Policy Table. This table basically contains rules that identify the DNS suffixes that must communicate through the corporate intranet to be reachable.

Although DirectAccess looks similar to a VPN in terms of creating a tunnel between the computer and the corporate environment, there's an important difference. Behind the scenes, there are actually two tunnels used with DirectAccess. The first tunnel is the infrastructure tunnel, which is established when the machine is turned on. It allows key IT infrastructure systems to talk to the machine and perform management. After a user logs on, a second tunnel is created. This intranet tunnel allows the user to access the systems on the corporate network. Because there's no user action required, the user can access corporate resources seamlessly. Note that it's possible to use only the infrastructure tunnel for management and not the intranet tunnel, in which case users wouldn't be able to access the corporate resources.

As you can see, DirectAccess is great for users, but it's even better for the IT department. With DirectAccess, all the communications traveling across the Internet are authenticated and encrypted using IPSec, which gives users a seamless but highly secure connection from their machines to the corporate network. In addition, because users' computers are connected to the key IT infrastructure systems whenever the users turn on their computers, it's easy to manage those computers.

DirectAccess is built on IPv6. Although the industry is certainly moving toward using IPv6, it's a very slow transition and many networks, including the Internet, are primarily using IPv4. As a result, DirectAccess leverages IPv6 transition technologies to establish communication between the DirectAccess server and those clients using IPv4 to connect to the Internet. The common transition technologies being used are Teredo and IP over HTTPS (IP-HTTPS), both of which allow the tunneling of IPv6 within IPv4. Technically, the 6to4 transition technology can also be used. However, 6to4 doesn't work when

the client is behind a Network Address Translation (NAT) device. Because most Internet-based clients are behind a NAT device of some kind, 6to4 is rarely used in real life.

DirectAccess is a fantastic technology, but it's highly unlikely that you'll be able to get rid of your VPN solution. For example, DirectAccess in Server 2012 works only with domain-joined Windows 8 Enterprise and Windows 7 Enterprise machines. For all other devices—e.g., home machines that aren't domain-joined, non-Enterprise editions of Windows 8 and 7, mobile phones, non-Windows devices—you'll still need to leverage a VPN. I recommend that when you can use DirectAccess, use it. For everything else, you'll need to continue to use a VPN. (If you're struggling to understand the difference between a VPN and DirectAccess, I've heard this summary: The VPN connects the user to the network, whereas DirectAccess extends the network to the computer and user.)

DirectAccess in Server 2012

To deploy a usable DirectAccess implementation in Server 2008 R2, you really need to use the Forefront UAG. It provides a simpler setup experience and enables support for IPv4 targets. Server 2012 includes all of Forefront UAG's technologies related to DirectAccess, such as DNS for IPv6 to IPv4 (DNS64) and Network Address Translation for IPv6 to IPv4 (NAT64). As a result, DirectAccess will work in a Server 2012 network, without requiring you to install an additional product like Forefront UAG.

Server 2012 introduces a new connectivity-related role named Remote Access. Through this role, you can manage DirectAccess and VPNs (including site-to-site VPNs) as a unified service.

Server 2012 also introduces multi-site support and geographical awareness, which means you can have multiple DirectAccess deployments (single servers or arrays) at different locations, and clients will use whichever site is closest based on response time. The response time is determined using an HTTP probe that tests connectivity to all

the DirectAccess deployments identified in its policy. If a site fails, clients will use the remaining DirectAccess deployments. Multi-site awareness is a Windows 8 native capability. However, if you really need this capability for Windows 7 clients, there are ways to make it work to a certain extent. (There are some constraints.)

If you have only Windows 8 clients, the requirement for a PKI is removed in Server 2012 DirectAccess. In Windows 7 clients, PKI is needed for IPsec, but now IPsec can actually use Kerberos tickets through the Kerberos proxy running on domain controllers (DCs).

One common concern when using DirectAccess is security. It's a completely automated technology—the computer automatically connects to the intranet through its secure tunnel and is always on. For authentication, DirectAccess uses:

- The computer certificate and computer account (using NTLM) to establish the infrastructure tunnel
- The computer certificate and user account to establish the intranet tunnel after a user logs on using Kerberos

You could consider this a kind of “1.5 factor” authentication because you have the certificate bound to the machine and user password. However, many organizations require two-factor authentication, so smart cards are often used for DirectAccess deployments that have Windows 7 clients. For Windows 8 clients, physical smart cards are no longer required (but are still supported) because you can use the new Windows 8 virtual smart card feature that leverages the machine's Trusted Platform Module (TPM). This makes deployments much simpler, reduces costs, and increases security. A one-time password is also supported with Server 2012 DirectAccess.

Changes in DirectAccess Requirements

Because Microsoft made many changes in the way DirectAccess works in Server 2012, there are many changes in what is required to use it. Here are the most noteworthy changes made to the IP address

requirements, Active Directory (AD) requirements, server requirements, and client requirements.

Changes in IP address requirements. In Server 2008 R2, two consecutive public IPv4 addresses are required for the DirectAccess service because Teredo requires two IP addresses to ascertain the type of NAT device the client is behind. (For more details, see “[Teredo: Tunneling IPv6 over UDP through Network Address Translations \(NATs\)](#).”) In Server 2012, you can place the DirectAccess server behind the NAT device. Thus, if the NAT device is present in the demilitarized zone (DMZ), the DirectAccess server can be deployed using private IP addresses. This means you can now set up DirectAccess on home computers if desired.

It’s important to note that while most DirectAccess deployments are behind the NAT device, the most optimal deployment over the Internet is using Teredo, which doesn’t support NAT. So, having two consecutive public IP addresses is ideal, but if it isn’t possible, you just need to use IP-HTTPS. (Most times, clients don’t have public IP addresses, so 6to4 can’t be used. Even when it can be used, such as on mobile networks, I’ve seen problems arise.)

Changes in AD requirements. A single DirectAccess deployment can now service multi-forest deployments, provided there’s a bi-directional trust between the forests. Although you still need to configure cross-forest scenarios, the graphical tools in DirectAccess do all the heavy lifting, provided that a bi-direction forest trust exists between the forest containing the DirectAccess servers and the forests containing the users.

Changes in server requirements. In Server 2012, Server Core is fully supported for the DirectAccess server. Interestingly, Microsoft now recommends running the DirectAccess server on a virtual machine (VM). If you run it on a VM, you can still offload the IPsec traffic to the network adapter (assuming it supports IPsec offload) and still fully leverage technologies such as Receive Side Scaling for best performance.

Previously, Microsoft recommended running DirectAccess on physical hosts because of the high workloads related to encryption. In Server 2008 R2, DirectAccess uses double encryption if HTTPS is used because encryption occurs for the IPsec traffic and then again for the HTTPS traffic. In Server 2012, DirectAccess can use IP-HTTPS NULL encryption instead of double encryption, thereby reducing the overhead for the client and DirectAccess server. Ultimately, this means less resource usage and more users per server—I've heard as high as four times the number of users. In Server 2012, using IP-HTTPS is almost on performance parity with using Teredo.

Changes in client requirements. When you use Windows 7 clients with DirectAccess in Server 2012 or Server 2008 R2, you need to install a separate DirectAccess Connectivity Assistant (DCA), which gives a system tray icon that shows the DirectAccess connection state. When you use Windows 8 clients with DirectAccess in Server 2012, the DCA isn't needed because DirectAccess support is integrated with the rest of the networking features in the Windows 8 OS. The DirectAccess connectivity status is shown alongside the status of other connections (e.g., wireless connections) in the networking UI. The networking UI also exposes a DirectAccess properties interface, which is typically used for troubleshooting. The DirectAccess properties interface allows DirectAccess logs to be exported to a file, which can then be sent to the IT Help desk. It also allows users to override the chosen DirectAccess site and specify which site they'd rather connect to.

Offsite provisioning is useful to set up clients that aren't physically connected to the corporate network. In Windows 8 clients, full offsite provisioning is possible, including a Windows-To-Go based installation. (Windows-To-Go isn't possible when using Windows 7 clients with Server 2012 DirectAccess.) It's still necessary to connect to a corporate network resource to download the DirectAccess setup file, but this could be as basic as a secure website. (Offsite DirectAccess configuration is possible in Server 2008 R2 but requires a time-intensive workaround that involves creating a temporary VPN connection.)

Setup and Management of DirectAccess in Server 2012

It's no exaggeration to say that deploying a DirectAccess server is quick and easy in Server 2012. For example, the following describes how to perform a complete single-server DirectAccess deployment for a small or midsized organization. This deployment assumes that an external DNS name exists and points to the correct IP address. Here are the steps:

1. Install the Remote Access role and its default options using the command:

```
Install-WindowsFeature RemoteAccess  
-IncludeManagementTools
```

2. Open the Remote Access Management Console.
3. Click the Getting Started Wizard option. This will launch an express wizard. (There's also an expert option available.)
4. On the Configure Remote Access page, select the *Deploy DirectAccess only* option.
5. On the Remote Access Server Setup page, you'll see your topology choices, which are based on the capabilities of your DirectAccess server. For example, if a server has only one network adapter, it can only use one type of topology. The types of topologies include Edge (where there are two network adapters—one for the Internet and one for the private network), Back (where the DirectAccess server is behind a NAT device, with connections to the DMZ and the private network), and Single (where there's one network adapter with a single network connection). After selecting the appropriate option for your topology, you need to specify a public IP address or an externally resolvable DNS name that clients will use to connect to your DirectAccess server.
6. When the page summarizing the settings is displayed, click OK. DirectAccess is now set up and ready to be used.

**Video**

John Savill
demonstrates how to
set up a DirectAccess
server in Windows
Server 2012



If you want to watch these steps being performed, check out the accompanying video.

Besides setting up the DirectAccess server, you need to set up the clients. When you deploy a DirectAccess server, DirectAccess automatically creates a GPO named DirectAccess Client Settings, which contains the DirectAccess client configurations. If you want a client to use DirectAccess, you need to join it to the domain and apply that GPO. This can be difficult if that client isn't connected to the corporate network, in which case you need to join the domain offline and apply the policy. To accomplish this in a Windows 8 client, you can use the Djoin.exe command-line utility.

To use Djoin.exe, you first need to provision the machine account in AD and save the provisioning data to a file. For example, the following command provisions a computer account named DAclient Example in the domain savilltech.net, specifies that the DirectAccess

Client Settings GPO be applied, and saves the provisioning information in the clientda.txt file:

```
Djoin.exe /provision  
/domain savilltech.net  
/machine DAclientExample  
/policynames "DirectAccess Client Settings"  
/savefile clientda.txt
```

(Although this command wraps here, you'd enter it all on one line in Cmd.exe. The same holds true for the next command.) After the client has created the clientda.txt file, you can run the following command to request an offline domain join the next time the client starts up:

```
djoin.exe /requestodj /loadfile clientda.txt  
/windowspath %windir% /localos
```

So, the next time the client starts, it will join the savilltech.net domain and apply the DirectAccess Client Settings GPO.

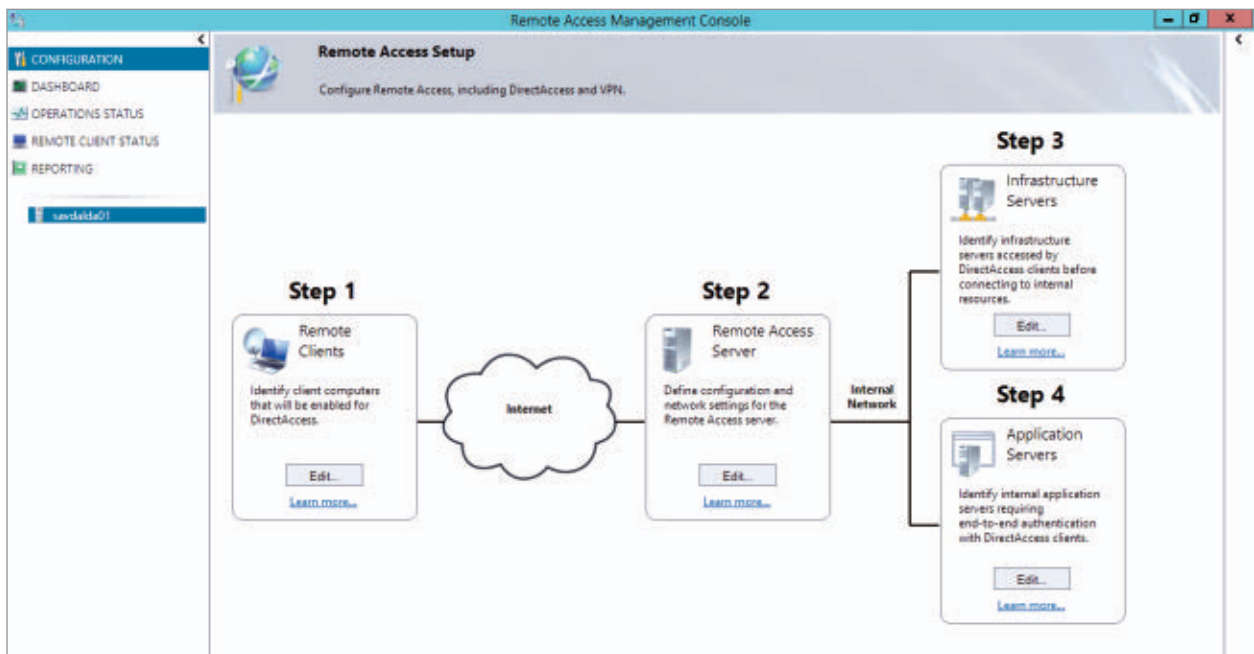
Note that if you have Windows 7 clients, there are some additional steps required. However, I won't cover them here, because my goal is to show you how easy it is to apply Server 2012 DirectAccess when you're using Windows 8 clients.

The DirectAccess server and client setups I demonstrated here are minimal deployments. Realistically, you'll probably want to make a few changes:

- You'll probably want to create a separate group that contains the computer accounts you want to enable for DirectAccess (e.g., a group named DA_Clients). The default is to use the domain's Domain Computers group, which will deploy DirectAccess to every machine in your domain.
- You might want the DirectAccess Client Settings GPO to be applied to specific groups rather than the root domain, which is

the default if you're logged on as a domain administrator when you run the commands to set up the DirectAccess service. Fortunately, by design, the GPO will be applied to the computer group specified in the Remote Access Management Console. This means that the application of DirectAccess will be controlled if you specify an alternate computer group instead of the Domain Computers group, even if the GPO is linked to the root domain.

These customizations are possible and desirable in most instances. You can easily make them (as well as many other customizations) after DirectAccess is deployed using the Remote Access Setup page shown in Figure 1, which is the starting page for DirectAccess configuration.



In addition to tweaking the initial deployment, you can use the Remote Access Setup page throughout the life cycle of DirectAccess, as your company's needs change. For example, you can use it to add servers, change public names, and change the topology. In addition,

Figure 1
Remote Access Setup
Page

there's full Windows PowerShell support for automating tasks such as applying GPO and configuration changes. In Server 2008 R2, you would need to get up at 2:00 A.M. to manually apply GPO and configuration changes through a GUI.

A Vast Improvement

As you can see, DirectAccess is far easier to set up in Server 2012 than it was in Server 2008 R2. Plus, it offers a great end-user experience and powerful management capabilities for your IT department—capabilities far beyond what's possible with a manually initiated VPN connection. But remember, you'll still need a VPN for those non-DirectAccess-capable machines. Fortunately, Server 2012 offers a great VPN experience as well. ■

PowerShell Basics: Select-Object

How to use this crucial cmdlet

Windows PowerShell is all about objects. That should be the mantra for every PowerShell-loving IT professional. PowerShell's elegance is derived in large part by how you work with objects in the pipeline. And one cmdlet that you absolutely need to master is `Select-Object`. This cmdlet does exactly what the name suggests: It selects objects. Actually, it can select objects in a few ways, which I'll explain and demonstrate. I'll be using PowerShell 3.0 for the demonstrations.

Selecting a Number of Objects

The first way to use the `Select-Object` cmdlet (which has the alias of `Select`) is to select the first or last X number of objects. This is especially useful when you only need a sampling or subset of data. You use the `-First` parameter to select from the beginning of the data and the `-Last` parameter to select from the end of the data.

For example, suppose you're interested in discovering more about the methods and properties of the eventlog object, which you'd like to use to manage your event logs. You can use a command like this:

```
Get-EventLog -List | Select -First 1 | Get-Member
```

Here's how this command works:

1. The `Get-EventLog` cmdlet gets a list of the eventlog objects.
2. The list is piped (i.e., sent) to the `Select-Object` cmdlet. The `-First 1` parameter tells PowerShell to select the first eventlog object in that list.



Jeffery Hicks

is a Windows PowerShell MVP with almost 20 years of IT experience. He works as an independent consultant, trainer, and author. His latest book, with Don Jones and Richard Siddaway, is *PowerShell in Depth: An administrator's guide* (Manning, 2013).



3. That object is piped to the `Get-Member` cmdlet, which lists its properties and methods. (If all I wanted was member information, I wouldn't need to use `Select-Object`. But I wanted to demonstrate how to use it. Plus, if the initial command returned numerous objects, this technique would improve performance.)

When it comes using the `Select-Object` cmdlet to select a number of objects, a more common usage is using it after some objects are sorted. For example, suppose you want to find the most recently modified file in your `Scripts` folder. To accomplish this, you can run the command:

```
Dir C:\scripts -File | Sort LastWriteTime | Select -Last 1
```

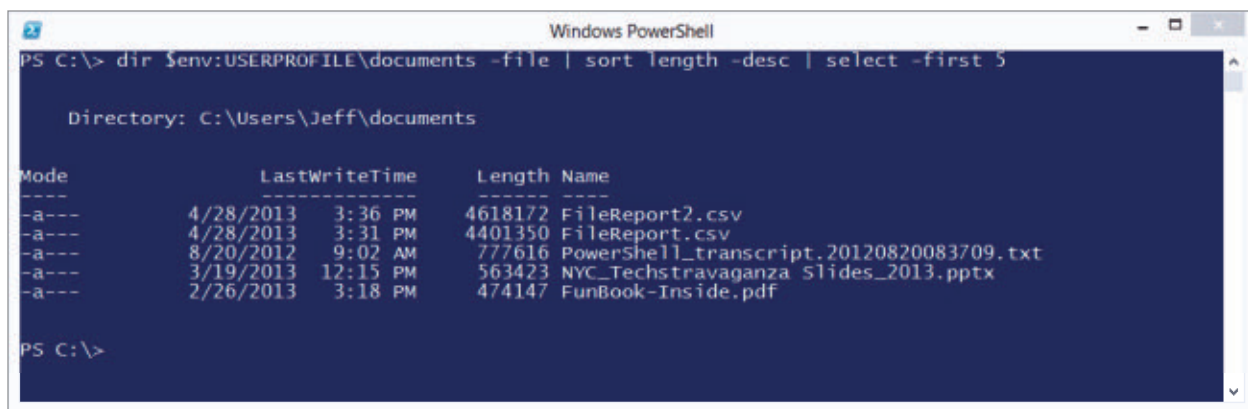
Here's how this command works:

1. The `Get-ChildItem` cmdlet (which has the alias *Dir*) with the `-File` parameter lists the files contained in the specified directory (`C:\scripts`). Note that the `-File` parameter was introduced in PowerShell 3.0. If you aren't running that version, you'll get an error.
2. The files are piped to the `Sort-Object` cmdlet (which has the alias *Sort*) and sorted based on the `LastWriteTime` property. The `-Property` parameter is the first positional parameter, so its parameter name (`-Property`) isn't required.
3. The sorted list is piped to the `Select-Object` cmdlet with the `-Last 1` parameter to select the last object (i.e., the most recently modified script).

The `Sort-Object` cmdlet's default sort order is in ascending order, but you can change that behavior by including the cmdlet's `-Descending` parameter. Here's a good example:

```
Dir $env:USERPROFILE\Documents -File |  
    Sort Length -Descending |  
    Select -First 5
```

(Although this line wraps here, you'd enter it all on one line in the PowerShell console. The same holds true for the other commands that wrap.) As Figure 1 shows, this command lists the five largest files in my Documents folder. To do so, it retrieves all the files, sorts them based on the Length property in descending order, and selects the first five files. The objects written to the pipeline are unchanged. All this command does is tell PowerShell how many objects to select, then passes those selected objects to the pipeline.



```

PS C:\> dir $env:USERPROFILE\documents -file | sort length -desc | select -first 5

Directory: C:\Users\Jeff\documents

Mode                LastWriteTime         Length Name
----                -
-a---            4/28/2013   3:36 PM      4618172 FileReport2.csv
-a---            4/28/2013   3:31 PM      4401350 FileReport.csv
-a---            8/20/2012   9:02 AM       777616 PowerShell_transcript.20120820083709.txt
-a---            3/19/2013  12:15 PM       563423 NYC_Techstravaganza_Slides_2013.pptx
-a---            2/26/2013   3:18 PM       474147 FunBook-Inside.pdf
  
```

Figure 1

Listing the Five Largest Files in the Documents Folder

You can go a step further and calculate the total size of the five largest files using the Measure-Object cmdlet with the -Sum parameter. As the following command shows, you pipe the five largest files to the Measure-Object cmdlet, telling it to sum the values of the Length property:

```

Dir $env:USERPROFILE\documents -File |
    Sort Length -Descending |
    Select -First 5 | Measure-Object Length -Sum
  
```

In my case, the command returned the sum of 10834708 bytes.

Selecting Properties

Many objects in PowerShell have more properties than you see by default, and you might need to see some of the nondefault properties.

This is another way to use `Select-Object`. You can select the properties you want to view. You use the `-Property` parameter and specify a comma-separated list of property names. For example, suppose you want to view the `DisplayName` and `Status` properties of the Background Intelligent Transfer Service (BITS) on your machine. You can use the `Get-Service` cmdlet to get the information for BITS and pipe the results to `Select-Object` with those properties specified, like this:

```
Get-Service bits | Select -Property DisplayName,Status
```

Figure 2 shows the results.

Figure 2
Selecting the
`DisplayName` and
`Status` Properties of
BITS

```
PS C:\> Get-Service bits | Select -Property DisplayName,Status
```

DisplayName	Status
-----	-----
Background Intelligent Transfer Service	Running

You can use any property or property set when you pipe something to `Select-Object`. The following command uses the `Get-Process` cmdlet to get information about the `Winword.exe` process and pipes it to `Select-Object`, which selects the `PSResources` property:

```
Get-Process Winword | Select PSResources
```

Like the `-Property` parameter of the `Sort-Object` cmdlet, the `-Property` parameter of the `Select-Object` cmdlet is positional, so you don't have to specify the parameter name (`-Property`). Figure 3 shows the results.

Figure 3
Selecting the
`PSResources` Property
of the `Winword.exe`
`ProcessS`

```
PS C:\> Get-Process Winword | Select PSResources
```

Name	: WINWORD
Id	: 3504
HandleCount	: 574
WorkingSet	: 92618752
PagedMemorySize	: 63451136
PrivateMemorySize	: 63451136
VirtualMemorySize	: 425492480
TotalProcessorTime	: 00:01:12.2440631

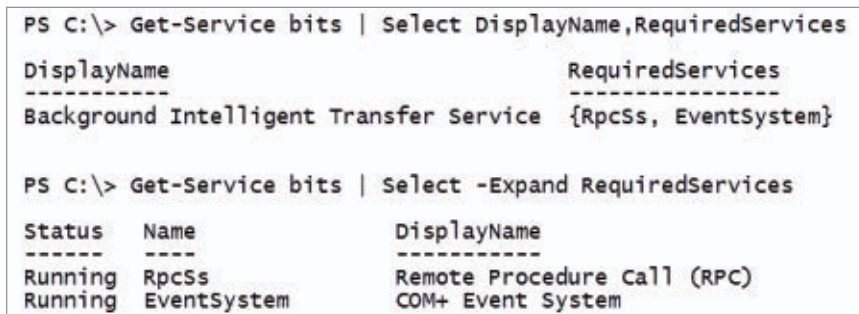
Selecting properties is the primary way most IT pros use Select-Object. You can also select properties using wildcards, which saves some typing.

Selecting Expanded Properties

Some object properties are collections of values or nested objects. When you select them, you might not get the output you expect. Take, for example, this command:

```
Get-Service bits | Select DisplayName,RequiredServices
```

In the output in Figure 4, notice that the value for RequiredServices is in curly brackets.



```
PS C:\> Get-Service bits | Select DisplayName,RequiredServices
```

DisplayName	RequiredServices
Background Intelligent Transfer Service	{RpcSs, EventSystem}


```
PS C:\> Get-Service bits | Select -Expand RequiredServices
```

Status	Name	DisplayName
Running	RpcSs	Remote Procedure Call (RPC)
Running	EventSystem	COM+ Event System

Figure 4

Selecting Properties That Are Collections of Values or Nested Objects

To see those values, you need to expand the property by using the -Expand parameter:

```
Get-Service bits | Select -Expand RequiredServices
```

Unfortunately, you can only expand a single property. The output of the second command in Figure 4 shows the expanded service objects that are required for BITS.

Selecting Isn't Formatting

When you select properties with the Select-Object cmdlet, sometimes the output isn't formatted very well. Figure 5 shows an example.

Figure 5

Selecting Isn't the Same as Formatting

```

Administrator: Windows PowerShell No Profile

PS C:\> get-process | select ID,name,WS,CPU

    Id Name                               WS                               CPU
    -- --                               -
    3508 audiodg                          7405568                        0.1872012
    1056 conhost                          7954432                        5.1168328
    416 csrss                             2822144                        1.4976096
    3620 csrss                             6852608                        7.8936506
    2656 dllhost                           1548288                        0.0468003
    2360 dpupdchk                           3371008                        0.0156001
    3312 dwm                               40574976                       139.776896
    2212 explorer                          78692352                       38.6258476
    0 Idle                                28672
    1996 ipoint                             12107776                       4.8828313
    2476 itype                              675840                        3.4008218
    3932 jused                                3395584                        0.0312002
    576 lsass                               7987200                        23.9305534
    1692 MsMpEng                           41762816                       540.543465
    1576 NitroPDFReaderDriverService3...  978944                        0.0312002
    1768 ONENOTEM                           831488                        0.0780005
    4052 powershell                        41914368                       8.112052
    3216 SearchIndexer                     22700032                       85.3169469
    568 services                           5181440                        14.1648908
    284 smss                                548864                        0.1872012
    2976 SnippingTool                       29589504                       4.5552292
    1392 spoolsv                             3768320                        0.4368028
    4060 SugarSync                          79753216                      2142.0341309
    676 svchost                             5054464                        4.8516311
  
```

Figure 6

Formatting the Output from the Select-Object Cmdlet

```

Administrator: Windows PowerShell No Profile

PS C:\> get-process | select ID,name,WS,CPU | format-table -AutoSize

    Id Name                               WS                               CPU
    -- --                               -
    3508 audiodg                          7352320                        0.1872012
    1056 conhost                          7962624                        5.5224354
    416 csrss                             2830336                        1.5288098
    3620 csrss                             6959104                        8.3460535
    2656 dllhost                           1548288                        0.0468003
    2360 dpupdchk                           3371008                        0.0156001
    3312 dwm                               39608320                      152.8497798
    2212 explorer                          77967360                      43.0874762
    0 Idle                                28672
    1996 ipoint                             12111872                       5.3508343
    2476 itype                              741376                        3.5412227
    3932 jused                                3395584                        0.0312002
    576 lsass                               8404992                       23.9461535
    1692 MsMpEng                           50274304                      540.8086667
    1576 NitroPDFReaderDriverService3...  978944                        0.0312002
    1768 ONENOTEM                           831488                        0.0780005
    4052 powershell                        42389504                       8.3772537
    3216 SearchIndexer                     23314432                       85.4261476
    568 services                           5226496                       14.1648908
    284 smss                                548864                        0.1872012
    2976 SnippingTool                       30138368                       6.3804409
    1392 spoolsv                             3768320                        0.4368028
    676 svchost                             5054464                        4.8672312
    728 svchost                             4538368                       6.9576446
  
```

This is to be expected. Selecting isn't the same as formatting. PowerShell does the best it can. If you want to pretty it up, you can pipe the

results from the `Select-Object` cmdlet to the `Format-Table` cmdlet, using the `-AutoSize` parameter, as I did in Figure 6.


Formatting must be at the end of your command, unless you are piping to one of the “Out” cmdlets such as `Out-File`. Most of the time, though, you’ll be selecting properties to obtain only the data you want so that you can do something with it, such as export to a .csv file or convert it to HTML.

Take the Time

Take the time to learn how to use `Select-Object`. You’ll be amazed by what you can accomplish in PowerShell when you use it. ■

CAN'T GET AWAY?

Get first-class education from your desk



Windows IT Pro offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to www.windowsitpro.com/events to see an up-to-date list of all online events.

Windows IT Pro

First-class education from the top experts in the industry.
Visit www.windowsitpro.com/events for a knowledge upgrade today!

Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.

Microsoft System Center 2012 SP1 Virtual Machine Manager User Roles

Having well-defined management permissions is important



Damir Dizdarevic

is manager of the Learning Center at Logosoft in Sarajevo, Bosnia and Herzegovina. He's an MVP for Windows Server Infrastructure Management, and an MCSE, MCTS, MCITP, and MCT.



Managing a private cloud environment typically involves much more than just providing templates that administrators can use to deploy new **virtual machines** (VMs). Administrators often need the ability to manage private clouds, VMs, and even virtual hosts. Thus, it's important to have well-defined management permissions in these **private cloud** environments.

Fortunately, Microsoft **System Center 2012 SP1** Virtual Machine Manager (VMM 2012 SP1) provides you with the ability to create user roles. With user roles, you can define a scope and the objects that administrators can manage. You can also define management operations that administrators can perform. After I describe the types of user roles available in VMM 2012 SP1, I'll show you how to assign administrators to existing user roles and how to create new user roles.

Understanding the User Roles

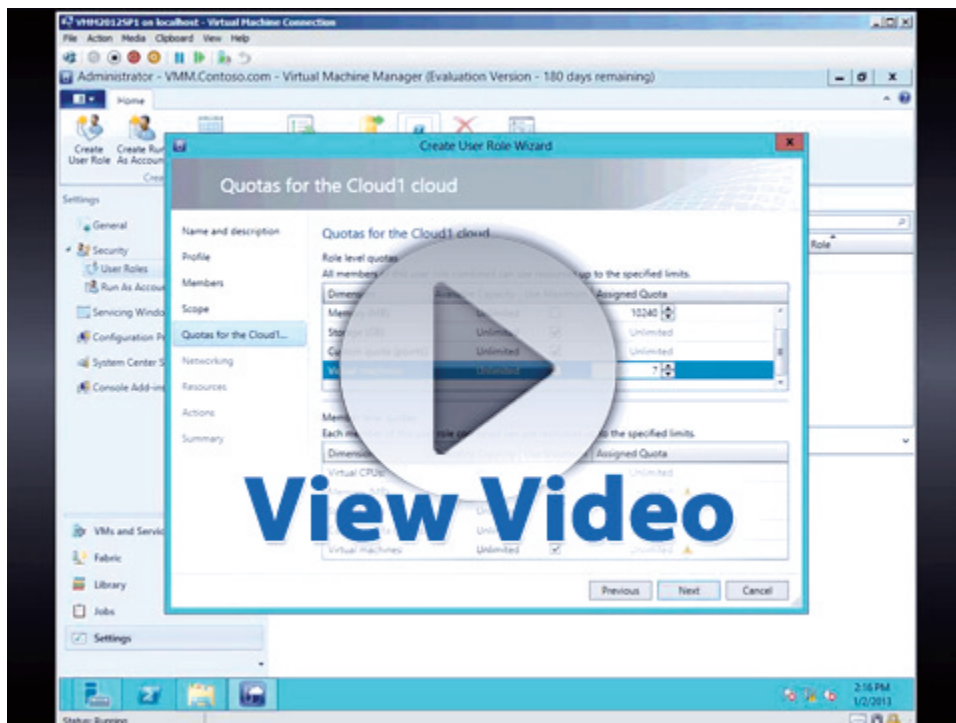
In VMM 2012 SP1, each user role you create comes with a set of permissions. Also, each user role is defined with a specific scope. In a private cloud environment, permissions are rarely delegated for the whole virtual infrastructure. Instead, permissions are delegated for "lower" levels, such as private clouds, host groups, or library resources.

VMM 2012 SP1 lets you define several types of user roles. The user roles that you can use include the following.



Video

Damir Dizdarevic
demonstrates System
Center Virtual Machine
Manager 2012
User Roles



Administrator. The Administrator user role comes predefined when you install VMM 2012 SP1. This default role has the widest management scope. Members of the Administrator user role can perform all administrative tasks on all objects (both virtual and physical) that VMM manages. Some tasks are specific to only this role and can't be delegated through any other role. For example, only members of the Administrator user role can add a standalone Citrix Systems XenServer to a VMM management server or add a Windows Server Update Services (WSUS) server for VMM fabric management. (Fabric is the term used to describe the infrastructure used to manage and deploy hosts, and to create and deploy VMs and services to a private cloud.) It isn't possible to redefine (i.e., narrow) the scope for the Administrator user role, so the number of members should be kept to a minimum. Typically, members are the administrators at the cloud provider company. The Administrator user role can create other

user roles and manage membership of any other user role. You should never assign users to this role.

Fabric Administrator. Members of the Fabric Administrator user role can perform all administrative tasks but within a defined scope. The scope can be a host group, private cloud, or one or more library servers. However, they can't modify any general VMM settings or modify the membership of the Administrator user role. If you want to give an administrator permission to fully manage a private cloud within VMM, this is the user role that you should use.

For hosted environments, this is very useful user role. For example, if you're a cloud provider and manage the virtual environment with VMM, you'll probably want to make your clients members of the Fabric Administrator user role so they can fully manage the objects and infrastructure within their private clouds. In this scenario, you'd define multiple user roles with the Fabric Administrator profile—one for each private cloud you create. Another scenario for using this role type is delegating other administrators with the ability to manage some portions of your virtual infrastructure. For example, you can give an administrator the right to manage specific host groups or library servers. Note that this user role is called the Delegated Administrator user role in the release to manufacturing (RTM) version of VMM 2012.

Read-Only Administrator. Members of the Read-Only Administrator user role can view but can't change the configuration settings for the VMM managed objects within a defined scope. They also can view the status of jobs executed within their management scope.

This user role is for auditing purposes. For example, if your virtual infrastructure is standardized and you want to make sure that change management is being properly managed, you can assign an auditing or change-management team member to this user role. You can also assign this user role to novice administrators who need to first familiarize themselves with the VMM configurations before being assigned to a user role with more permissions.

Tenant Administrator. This user role is specific to VMM 2012 SP1 and can't be created in VMM 2012 RTM. Members of the Tenant Administrator user role can define the scope of tasks performed by self-service users on their VMs, including creating and applying quotas on available resources. So, this is the user role you should use if you want to give an administrator permission to manage self-service users and the resources they consume.

Members of the Tenant Administrator user role can also manage VM networks, including managing and deploying their own VMs within a defined scope. The scope is limited to private cloud objects.

Application Administrator. Members of the Application Administrator user role can deploy and manage their own VMs within the scope and quotas defined by higher-level administrators. Note that this user role is called the *Self-Service User* user role in VMM 2012 RTM.

Assigning User Roles

Assigning an administrator to a user role in VMM 2012 SP1 is a pretty simple task. For example, if you want to add someone to the Administrator user role, you follow these steps:

1. Navigate to Settings in the VMM 2012 console, expand Security, and click *User roles*.
2. Double-click Administrator in the right pane.
3. Select the Members tab. Here you can add any user account from the Active Directory (AD) domain to which the VMM server belongs.

Note that you must use the VMM console or PowerShell to add an AD user account. You can't manage user roles from any AD utility.

Creating User Roles

You use the Create User Role Wizard to create new user roles. To open this wizard, you can navigate to Settings in the VMM 2012 console, expand Security, select *User roles*, and click the Create User Role

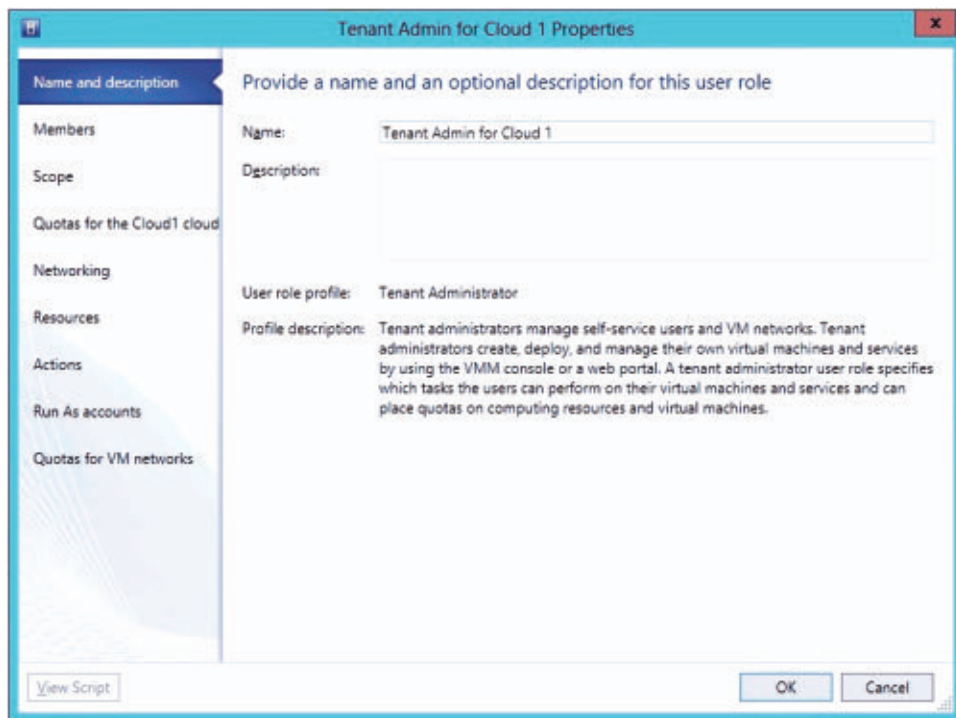
button. Alternatively, you can navigate to the Tenants node in the *VMs and Services* task pane, right-click the Tenants node, and select Create User Role. (Note that if you're using VMM 2012 RTM, you can't use the alternative method.)

The information that you need to provide in the Create User Role Wizard varies depending on the type of user role you're creating. For this reason, I'll describe the pages in the wizard rather than walk you through an example of how to create a particular user role.

Name and description. On this page, you need to provide the name and description of the user role, as shown in Figure 1. You should try to be as descriptive as possible, especially if you plan to create many user roles.

Profile. On this page, you choose the type of user role to create. As Figure 2 shows, the profiles from which you can choose are Fabric Administrator, Read-Only Administrator, Tenant Administrator, and

Figure 1
Providing the Name
and Description of the
User Role



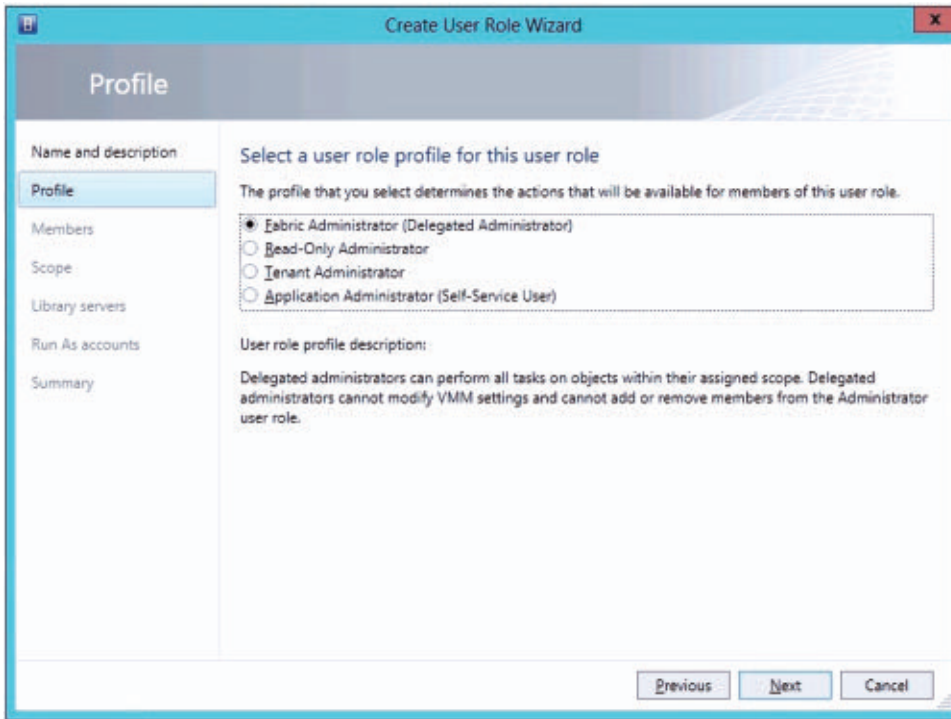


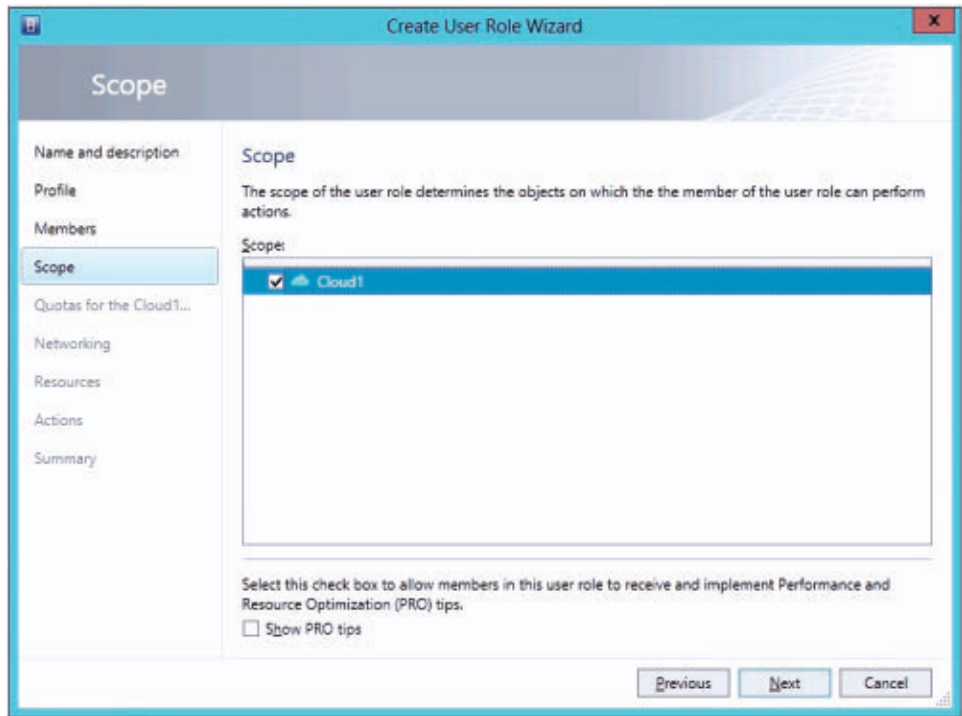
Figure 2
Choosing the Type of
User Role to Create

Application Administrator. The list doesn't include the Administrator user role because it comes predefined when you install VMM 2012, as mentioned previously.

Members. On this page, you can add user role members from AD. It isn't mandatory to do this when you're creating the user role. You can do it at any time by double-clicking the user role and navigating to the Members tab, as described in the "[Assigning User Roles](#)" section.

Scope. Figure 3 shows the Scope page, where you define the scope of the user role. This is a very important step. You need to select the VMM resources for which you want to give the user role permissions. If you're creating a Fabric Administrator or Read-Only Administrator user role, the available hosts groups and private clouds will be displayed. If you're creating a Tenant Administrator or Application Administrator user role, you'll see only the available private cloud objects. Be careful when selecting the resources. If

Figure 3
Defining the Scope of
the User Role



you make a mistake here, you can inadvertently provide access to the wrong resources.

Quotas for the cloud. This page is visible only if you're creating a Tenant Administrator or Application Administrator user role. As Figure 4 shows, you can define quotas for the private cloud objects that you've chosen on the Scope page. Defining quotas to limit resource usage is highly recommended. For example, you can define how many VMs each member of the user role can create and how much RAM can be used. Besides using quotas to limit resource usage, you can use them to monitor usage to determine whether you might need to add resources to your virtual environment.

Quotas are defined on two levels. You can define a total quota for a user role. You can also define quotas for each member of that user role. You can combine these two quota types so you have one general quota for the user role and specific quotas for each administrator who is a

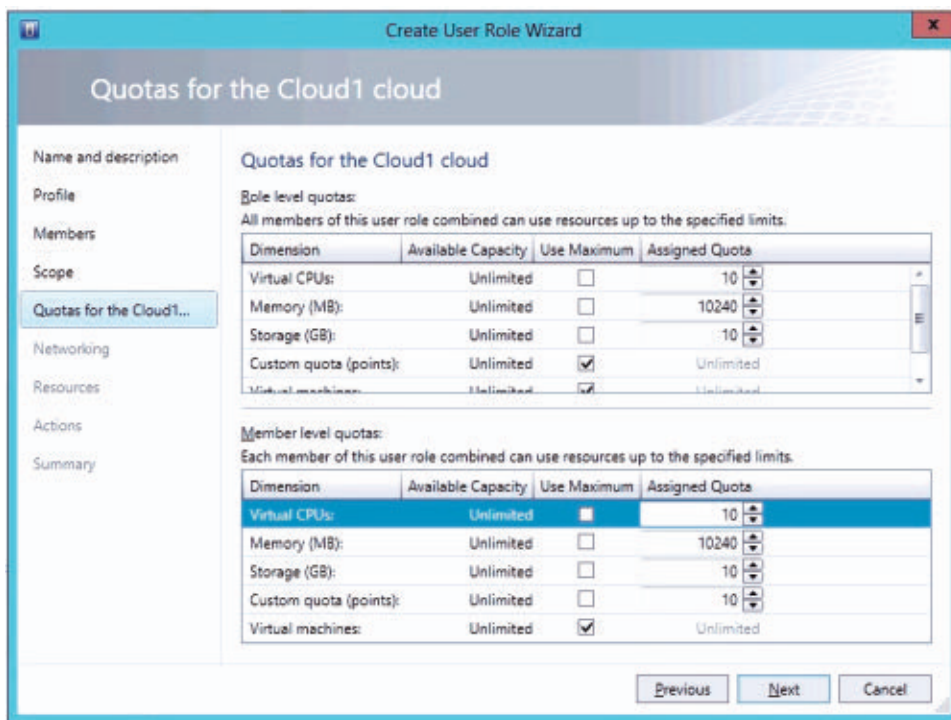


Figure 4
Defining Quotas for a
Private Cloud

member of that user role. When assigning quotas, make sure that you consider any quotas assigned to other user roles (if you have them). The system won't warn you if you oversubscribe, so make sure you don't.

Networking. Specific to only the Tenant Administrator and Application Administrator user roles, this page gives you the option to choose one or more VM networks that will be made available for usage. You also have the option to create new VM networks from this page.

Library servers. This page is visible only if you're creating a Fabric Administrator or Read-Only Administrator user role. In most environments, only one library server exists, so there will be no real choice. If multiple library servers are deployed, they usually host different resources. If you have more than one library server, you need to make sure you select the one that hosts the resources needed by the user role you're creating. In some scenarios, you can also deploy dedicated library servers for each private cloud you create.

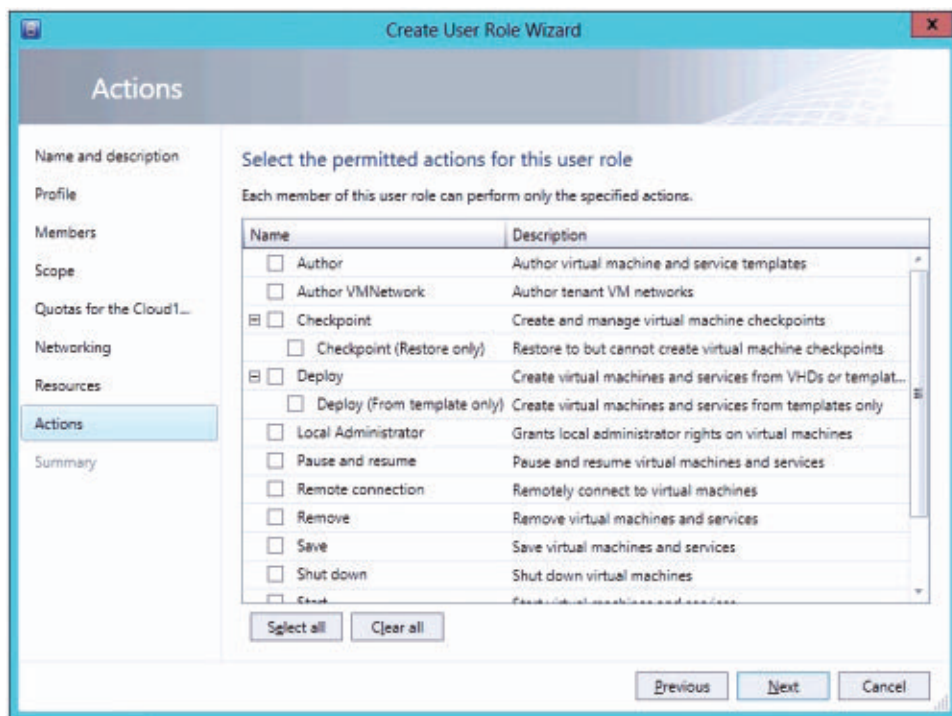
Resources. If you're creating a Tenant Administrator or Application Administrator user role, you need to choose specific resources from the library on the Resources page. It's important that you select the correct resources, especially if the administrators will be creating new VMs. You also need to define the data path for the data that the administrators will upload.

Actions. For the Tenant Administrator or Application Administrator user role, you'll have the option to choose specific actions that will be permitted. As Figure 5 shows, you can select actions such as Checkpoint (administrators can create and manage VM checkpoints) and Deploy (administrators can create VMs and services). Make sure that you understand the purpose of each action, taking into consideration the scope of the user role.

Run As account. This page appears if you selected the Author action on the Actions page for any of the user role types. On it, you can select

Figure 5

Selecting the Actions
That Will Be Permitted



a Run As account to be used by the members of the user role when executing tasks within VMM. A Run As account is a container for a set of stored credentials for a specific user account. You can create a Run As account before you run the wizard or when you run it.

Quotas for VM networks. This page appears if you're creating a Tenant Administrator user role and you selected the Author VMNetwork action on the Actions page. On it, you can define how many virtual networks each member of the Tenant Administrator user role can create or the total number of virtual networks that can be created by all the members of this user role.

Summary. On this page, you can review the settings you've entered before creating the user role.

An Important Part of Private Cloud Management

Creating and managing user roles is an important part of private cloud management. You should take care when configuring this aspect of VMM security, especially if you're working for a hosting provider that hosts private cloud environments for other companies. Using user roles is also a good way to control resource usage between various cloud administrators. ■

Windows Server 2012: Implement Continuously Available File Shares

Increase the range of storage options
for your mission-critical apps



**Michael
Otey**

is senior technical director for
Windows IT Pro and
SQL Server Pro.

Email



Continuously Available File Shares (CAFS) is an important new technology in [Windows Server 2012](#). At its basic level, Server 2012's CAFS feature takes Windows file sharing capabilities and scales them using a Server 2012 cluster. CAFS takes advantage of new Server Message Block (SMB) 3.0 capabilities to increase the availability of Windows Server file shares used for document storage and application support. Some of the new SMB 3.0 features that enable CAFS include SMB Scale-Out, SMB Direct, and SMB Multichannel.

The CAFS feature addresses problems that occurred in earlier implementations of highly available file servers on Windows Server failover clusters. Previous implementations provided high availability for file shares but were hampered by brief periods of downtime and a momentary loss of connectivity in the event of a failover. Such brief outages were usually acceptable for Microsoft Office-type applications that perform frequent file opens and closes, because these apps could reconnect and save changes after the failover completed. However, these same outages weren't acceptable for applications like Hyper-V or SQL Server, which hold files open for extended periods of time, and outages would result in data loss. Before the advent of Server 2012, Microsoft didn't support these types of server installations on file shares. Providing application support was one of Microsoft's primary design points for CAFS. While you can use CAFS



Video

Michael Otey demonstrates Windows Server 2012's Continuously Available File Shares (CAFS) feature

for simple client file sharing, CAFS is really targeted at supporting server applications. CAFS gives you the ability to take advantage of Windows Server's low-cost storage capabilities for mission-critical applications. CAFS provides continuous access to file shares with almost zero downtime.

Choose an Implementation

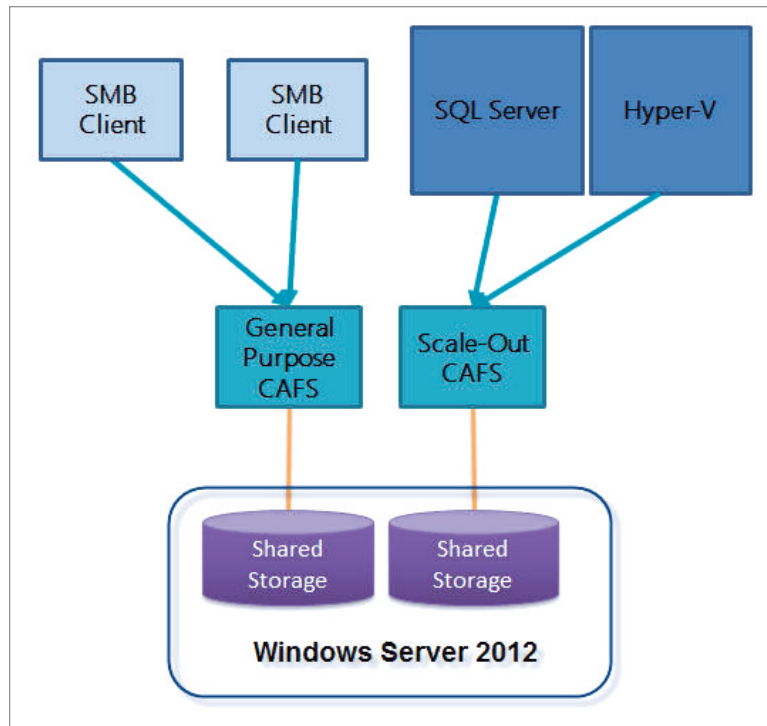
There are essentially two ways to implement CAFS:

- **General Purpose File Server**—Very much like the highly available file server support in Windows Server 2008 R2, the CAFS general use file server implementation allows a file share to be supported on a failover cluster. CAFS improves the availability and performance of this implementation with the new higher performance SMB 3.0 client access.
- **Scale-Out File Server**—The scale-out file server implementation is the new CAFS option for supporting applications like Hyper-V and

SQL Server with no downtime. This implementation is limited to four servers.

You can see an overview of the CAFS architecture in Figure 1.

Figure 1
Overview of
Continuously
Available File Shares
Architecture



One of the key technologies that enable CAFS is Server 2012's support for SMB Transparent Failover. SMB Transparent Failover lets file server services fail over to a backup node in the cluster so that applications with open files on the file server won't see an interruption in connectivity. CAFS addresses both planned maintenance and unplanned failures with zero application downtime.

Meet the Requirements

Because CAFS uses the SMB 3.0 features in Server 2012, the Server 2012 operating system is a definite requirement. CAFS is supported on both

the Server 2012 Standard and Server 2012 Datacenter editions. CAFS is not supported on the Essentials or Foundation editions.

In addition, CAFS requires a Server 2012 failover cluster. This means you must have a minimum of a two-node Server 2012 cluster. Server 2012 failover clusters support a maximum of 64 nodes. You can find step-by-step instructions on setting a Server 2012 failover cluster in my article [“Windows Server 2012: Building a Two-Node Failover Cluster.”](#) You also can watch a short [video in which I describe the process of building a two-node Server 2012 Failover Cluster.](#)

In addition to the cluster itself, the file server role must be installed on all cluster nodes. The clustered file server must be configured with one or more file shares that use the new continuously available setting. I provide more details about creating and configuring continuously available file shares later in this article.

For a two-node failover cluster, the cluster storage requires a minimum of two separate volumes (LUNs). One volume stores the shared files. This volume should be configured as a cluster shared volume (CSV). The other volume will function as the cluster witness disk. Most implementations use more volumes.

It’s also recommended that you design your network so there are multiple pathways between nodes. This prevents the network from becoming a single point of failure. Using network adapter teaming and multiple switches and/or redundant routers can add resiliency to your network configuration.

Finally, the SMB client computers must be running Windows 8 client or Server 2012 to take advantage of the new SMB Transparent Failover capability. When an SMB 3.0 client connects to a CAFS, the SMB client notifies the witness service on the cluster. The cluster picks a node to be the witness for this SMB connection. The witness node is responsible for switching the client to the new host in the case of an interruption of service, without requiring the client to wait for TCP timeouts.

Create a General Purpose CAFS

To configure a CAFS, open the Failover Cluster Manager on any of the nodes in the cluster. Then click the Roles node in the navigation pane. This displays existing roles in the Roles pane, as shown in the center of Figure 2.

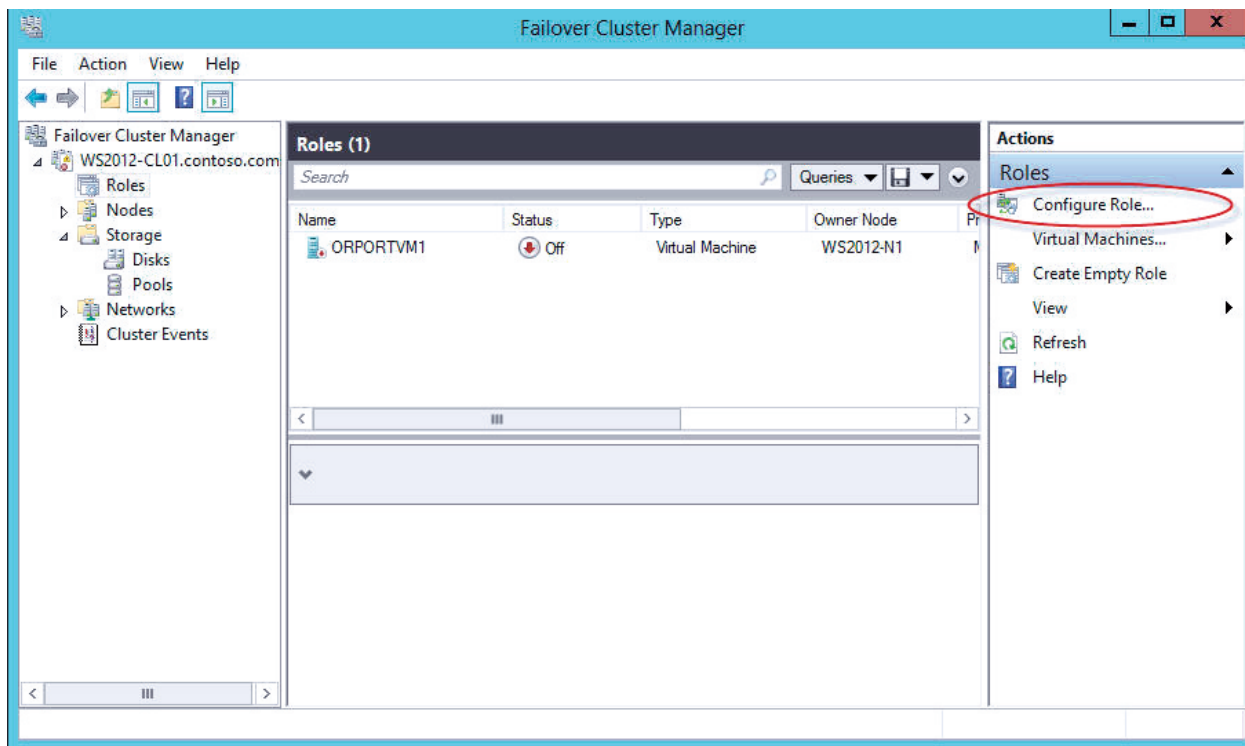


Figure 2
Failover Cluster
Manager

The cluster can support multiple roles and provide high availability capabilities to all of them. Figure 2 shows an existing, highly available virtual machine (VM). To create a new general purpose CAFS, click the *Configure Role* link highlighted in the Actions pane. This starts the High Availability Wizard shown in Figure 3.

Scroll through the list of roles until you see the file server role. The file server role supports both the general purpose and scale-out application types of CAFS. Select File Server and click Next to select the type of CAFS, which is displayed on the next screen, as Figure 4 shows.

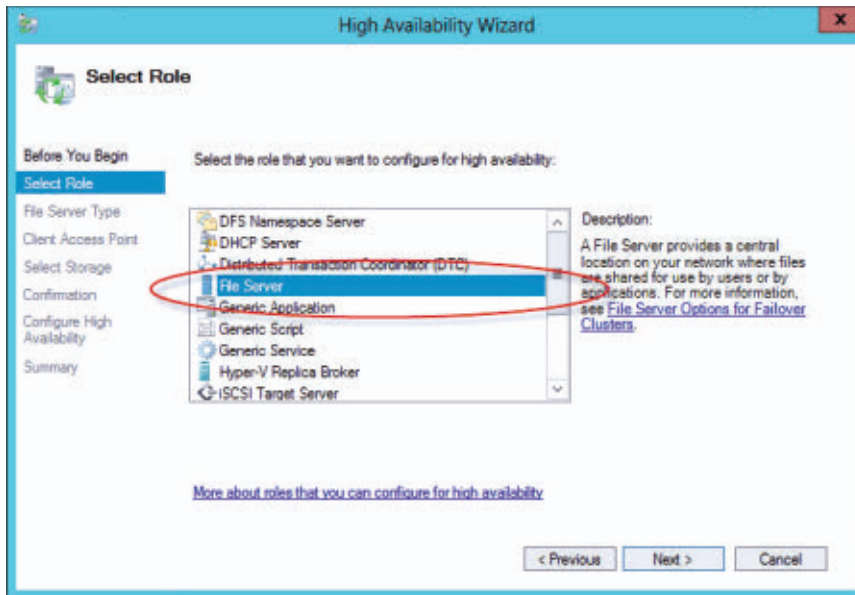


Figure 3
Adding the
File Server Role

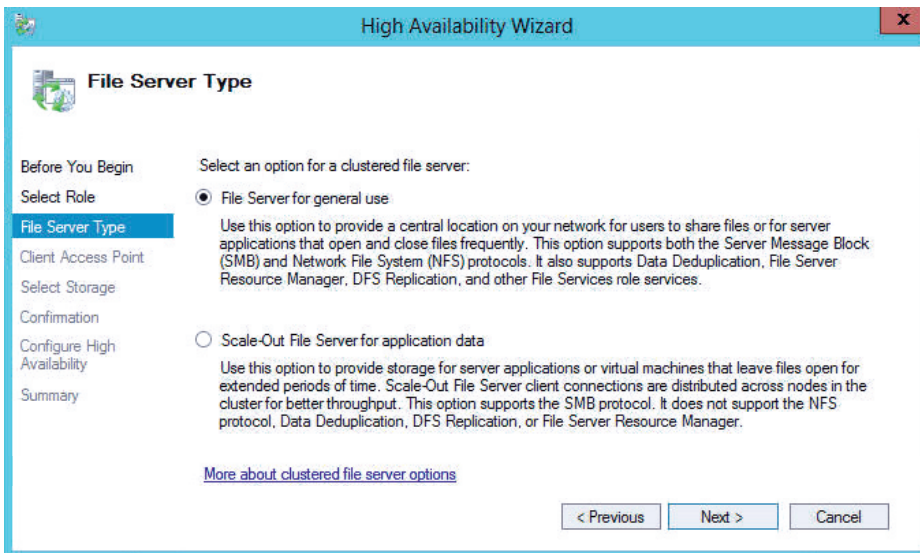
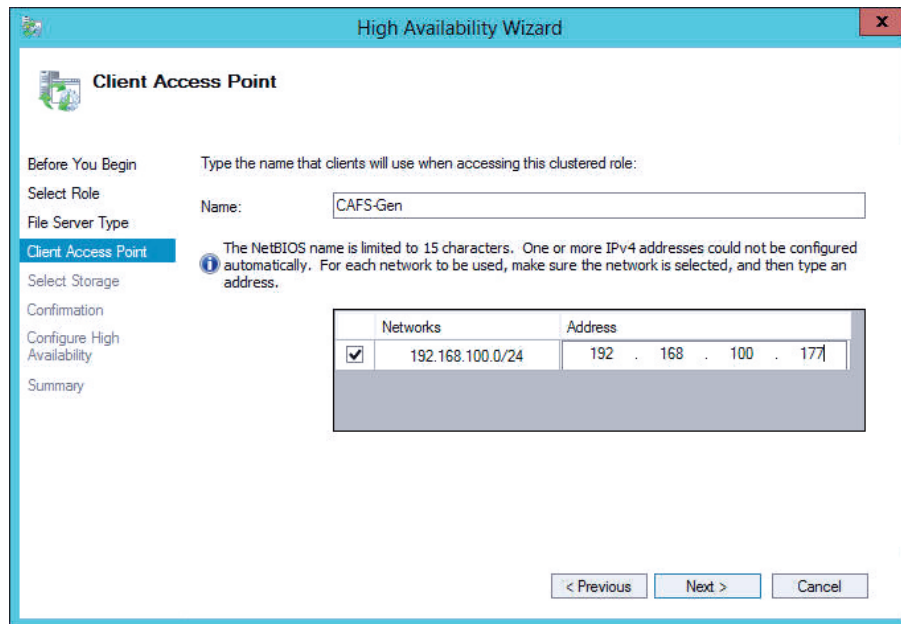


Figure 4
Selecting the File
Server Type to Create a
General Purpose
File Server

The File Server Type dialog box lets you choose between creating a *File Server for general use* or a *Scale-Out File Server for application data*. The general use option can be used for both Windows SMB-based file shares and NFS-based file shares. The general purpose CAFS

also supports data deduplication, DFS replication, and data encryption. Click Next to continue creating the general purpose CAFS. This displays the Client Access Point dialog box that Figure 5 shows.

Figure 5
Client Access Point for
General Purpose File
Server



To create a new general purpose CAFS, you must provide a server name that clients will use when they access the CAFS. This name will be registered in your DNS, and clients will use it like a server name. In addition, the general purpose CAFS also needs an IP address. In Figure 5 I named the service CAFS-Gen (for general purpose CAFS) and gave it a static IP address of 192.168.100.177. Clicking Next lets you select the cluster storage for the CAFS.

The Select Storage dialog box that Figure 6 shows lets you select the storage for the general purpose CAFS. The storage must be available to the cluster. In other words, it must be listed under the cluster's storage node and designated as available storage. You cannot use preassigned CSVs for your general purpose CAFS.

There are three disks that I could have used for this example, and I selected Cluster Disk 5 because I had previously allocated this storage

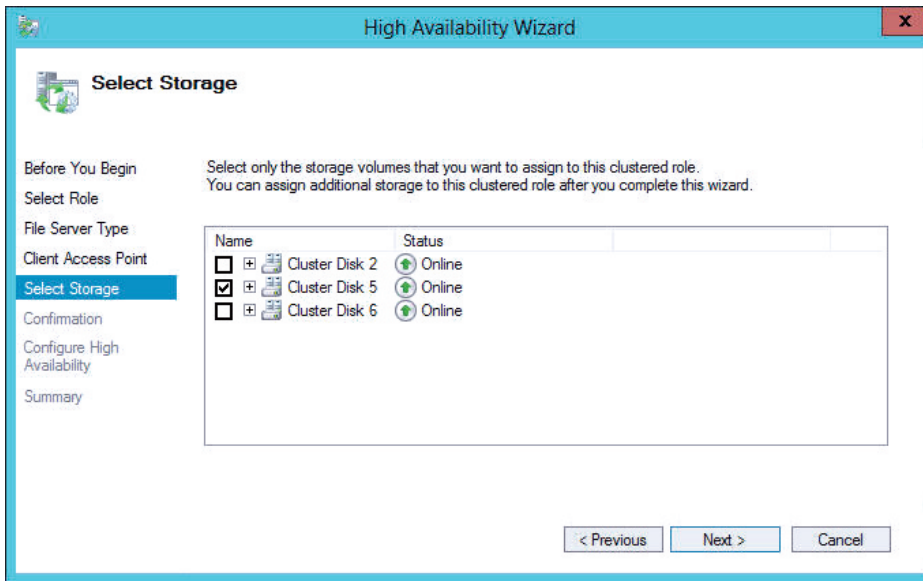


Figure 6
The Select Storage
Dialog Box

to the CAFS (Figure 6). However, you can select any of the available cluster disks. Clicking Next displays the Confirmation screen. At this point you can either confirm your selections or go back through the High Availability Wizard dialog boxes and make changes. If everything is OK, clicking Next on the Confirmation screen displays the Configure High Availability dialog box, which shows the progress of the CAFS configuration process. When it's complete, a Summary screen is displayed. Clicking Finish on the Summary screen closes the High Availability Wizard and returns you to the Failover Cluster Manager.

After creating the CAFS role, the next step is to create a continuously available file share that uses that role. Figure 7 shows that the CAFS-Gen role is actively running and that it uses the file server role. To add a new continuously available file share, select the *Add File Share* link in the Actions pane that you see on the right side of Figure 7. This displays a Task Progress dialog box that shows the progress of retrieving server information. Upon completion, the New Share Wizard displays.

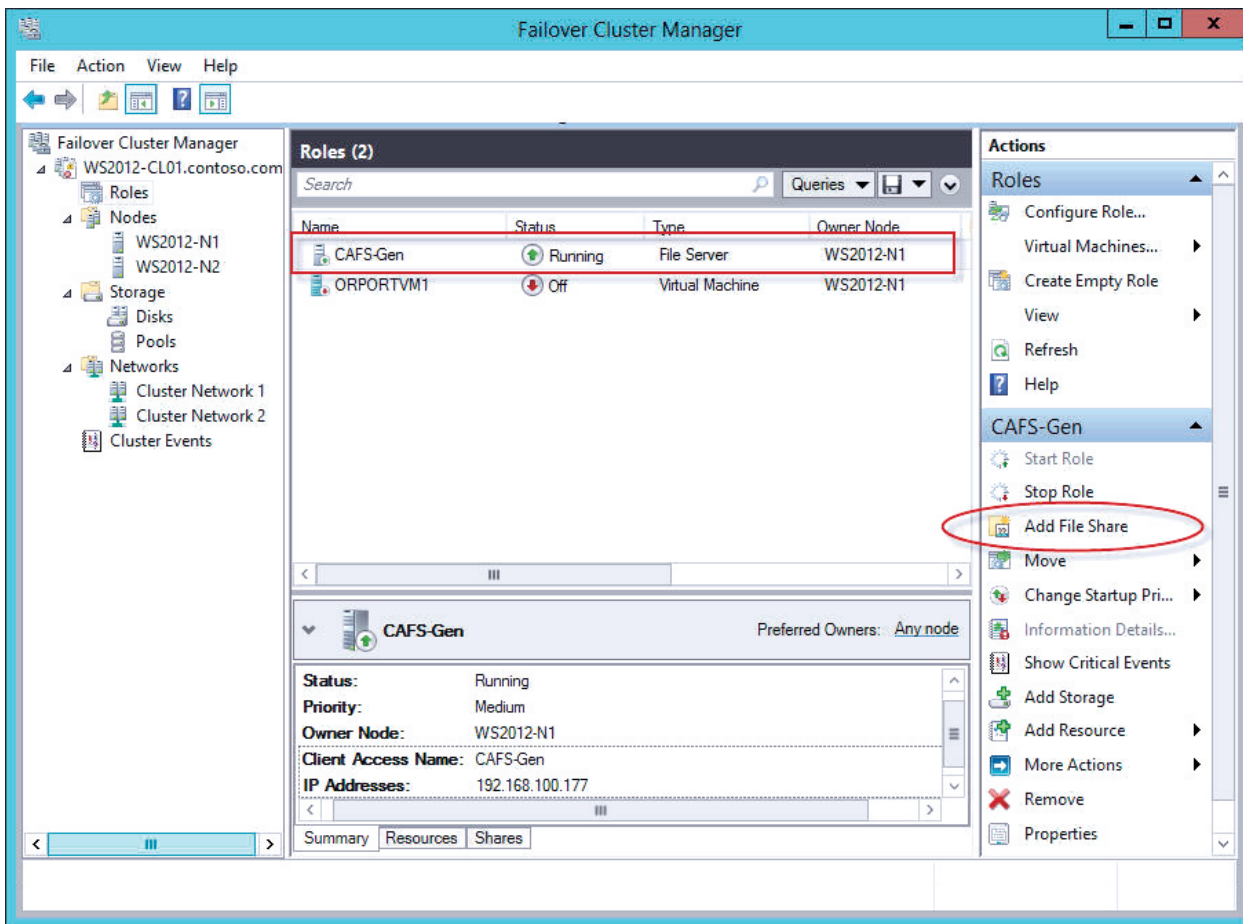
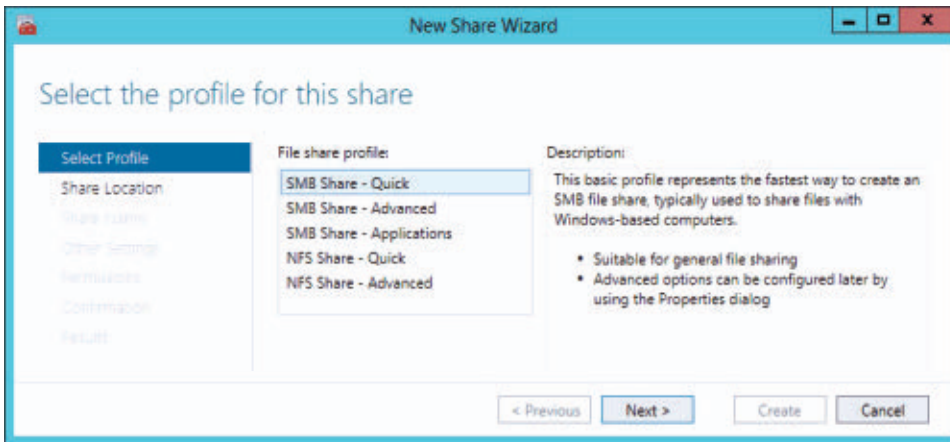
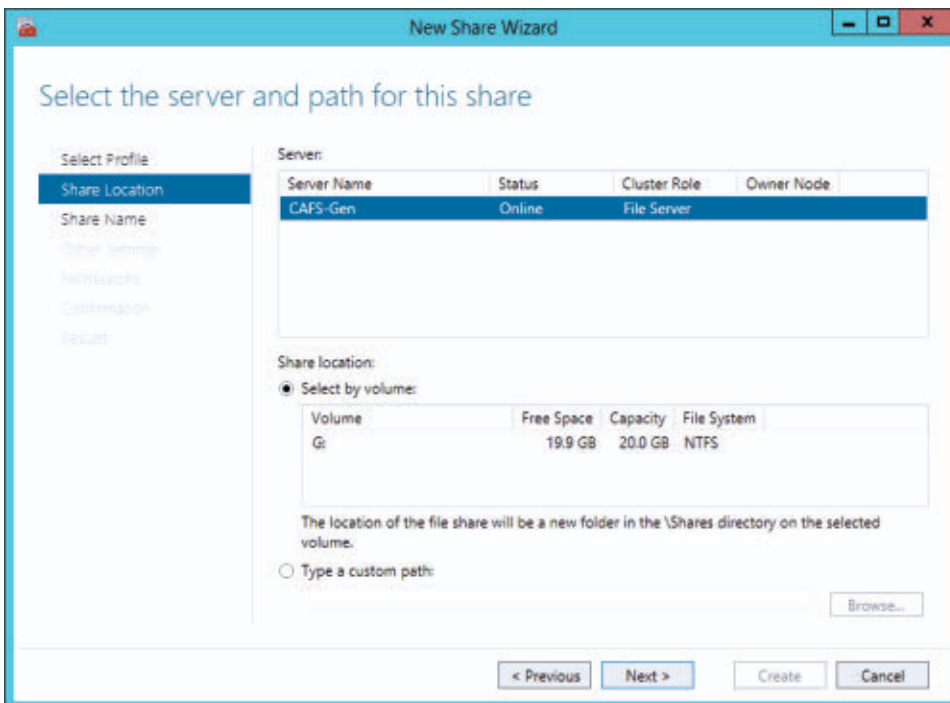


Figure 7
Adding a File Share

The New Share Wizard begins by asking what type of CAFS you want to create. You can choose to create either SMB or NFS types of CAFS. The *SMB Share—Quick* option creates a general purpose CAFS. The *SMB Share—Applications* option creates a highly available application share for applications like Hyper-V or SQL Server. I cover how to create a scale-out CAFS for applications later in this article. To create a general purpose CAFS, select the *SMB Share—Quick* option at the top of the list, as Figure 8 shows, and then click Next. The New Share Wizard displays the Share Location dialog box that Figure 9 shows.

**Figure 8**

Selecting a Profile for a General Purpose File Server

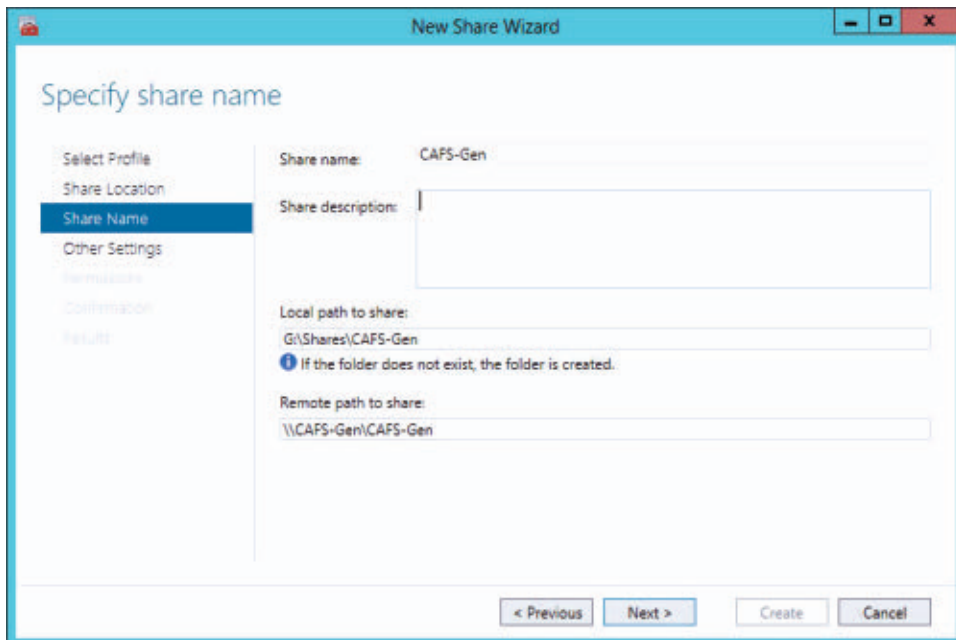
**Figure 9**

Share Location for General Purpose File Server

The name of the CAFS role is displayed in the Server Name box. Figure 9 shows the name of the CAFS-Gen role that I created earlier with a status of Online. You can select the location of the share using the options in the bottom half of the screen. In this example

the G drive was selected by default (see Figure 9). If you want to use a different drive, you can manually enter the alternative path in the *Type a custom path* text box at the bottom of the screen. In this example I stuck with the default G drive and clicked Next to display the Share Name dialog box shown in Figure 10.

Figure 10
Share Name for
General Purpose File
Server



The Share Name dialog box lets you provide a name for the file share. For simplicity, I used the same name for the general purpose CAFS that I used for the service: CAFS-Gen, but that isn't necessary. You can name the share any valid SMB name. In the center of the screen you also can see the local and remote paths to the CAFS. The local path for this example is G:\Shares\CAFS-Gen. The share will be accessed by networked systems using the path \\CAFS-gen\CAFS-Gen. Clicking Next displays the *Configure share settings* dialog box shown in Figure 11.

The *Configure share settings* dialog box lets you control how the share will be treated by the server. The *Enable continuous availability* check box is required to make the file share continuously available.

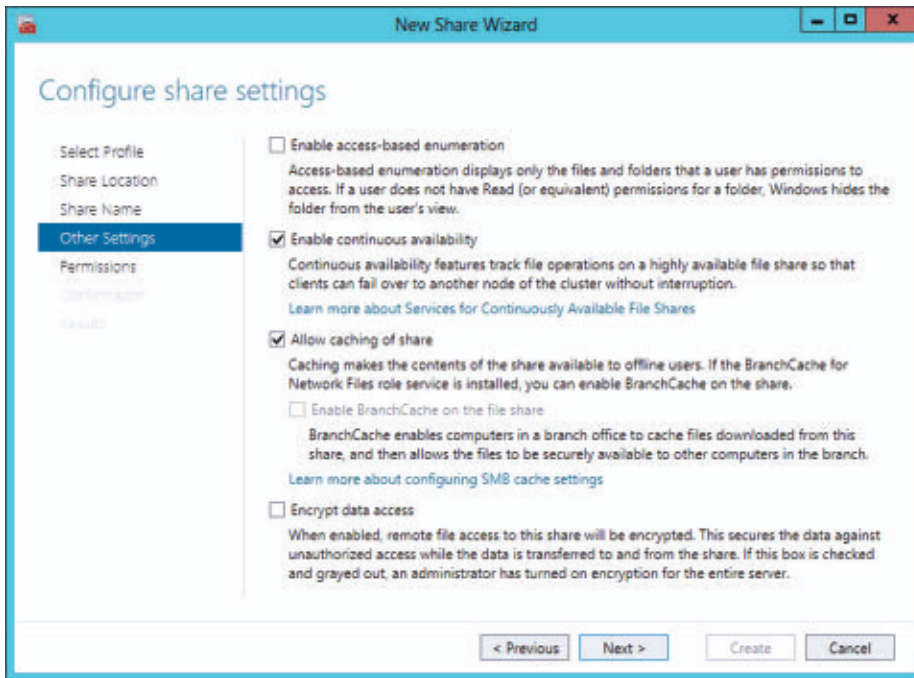


Figure 11
Configuring Share
Settings for the
General Purpose File
Server

This setting is checked by default. The *Enable access-based enumeration* setting controls whether users without permissions can view files and folders. This setting isn't checked by default. The *Allow caching of share* setting enables the contents of the share to be available to offline users via BranchCache. Finally, the *Encrypt data access* setting secures remote file access by encrypting the data transferred to and from the share. This setting is unchecked by default. Clicking Next displays the Permissions dialog box shown in Figure 12.

By default, the CAFS is created with Full Control given to the Everyone group. You'll probably want to change this for most implementations. I accepted the default permissions in this example. Clicking Next displays the Confirmation dialog box where you can view a summary of the choices you made in the previous New Share Wizard dialog boxes. You can click Previous to go back and change any of the settings. Clicking Create on the Confirmation dialog box creates the CAFS and sets the permissions for the share. After the CAFS share has been

Figure 12
Specifying Permissions
for the General
Purpose File Server

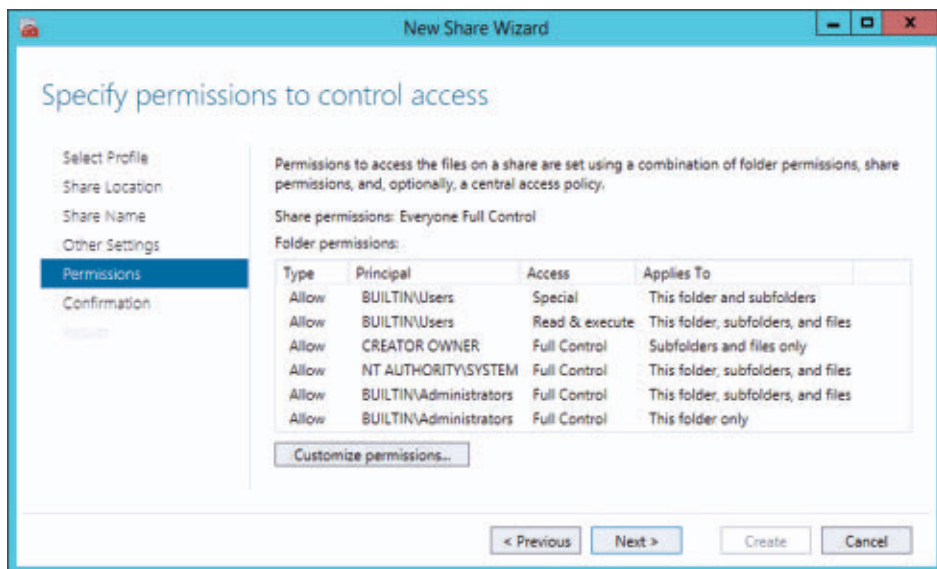
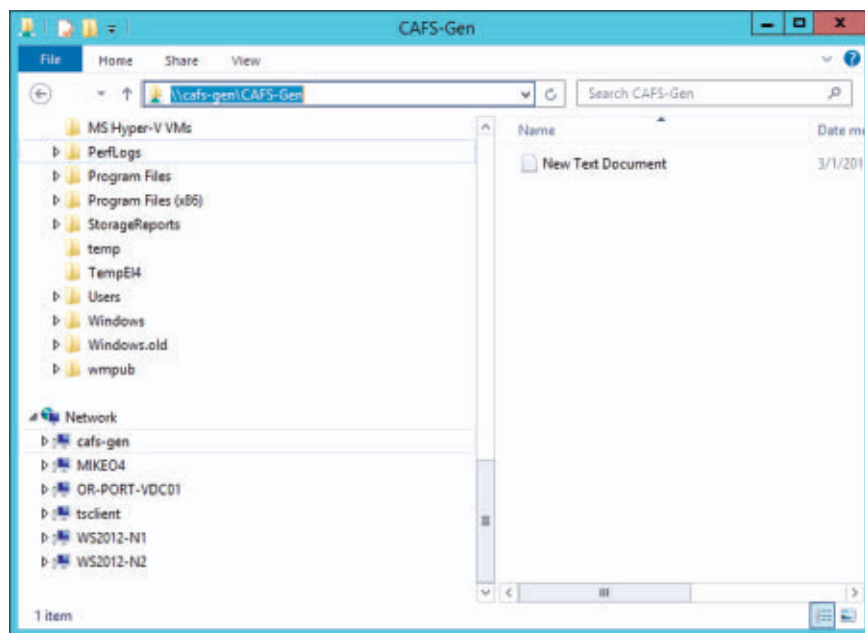


Figure 13
Accessing the CAFS by
Its Network Path



created you can access it like any other file share. Figure 13 demonstrates how to access the share by entering the \\cafs-gen\CAFS-Gen server and share name into Windows Explorer. At this point you can

populate the share with documents or other types of files that would benefit from the availability of a CAFS.

Create a Scale-Out CAFS

The primary purpose behind CAFS is to provide high availability to applications that store data on file shares. In the past, Microsoft didn't provide this kind of support for applications such as SQL Server that store their database on file shares. That changed with the release of Server 2012 and its support for the CAFS feature. Scale-out CAFS is implemented differently than general purpose CAFS. However, you use the same High Availability Wizard to create the scale-out option. To create a new CAFS for scale-out application support, select the *Configure Role* link in the Actions pane of the Failover Cluster Manager as demonstrated in Figure 2. Then on the Select Role dialog box, select the File Server role as shown in Figure 3. These first two steps are the same as for creating a general purpose CAFS. However, on the File Server Type dialog box, select the *Scale-Out File Server for application data* option as shown in Figure 14.

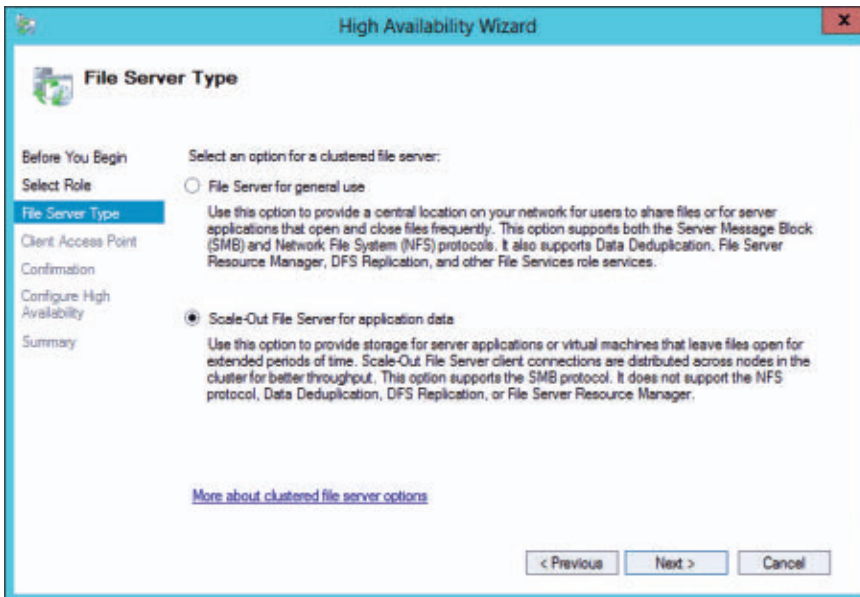
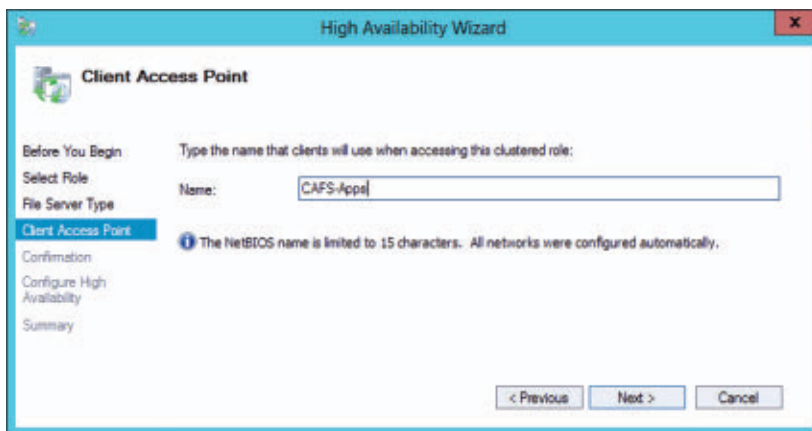


Figure 14
Selecting the File
Server Type to Create a
Scale-Out File Server

The scale-out file server option is designed for applications that leave their files open for extended periods of time. Clicking Next displays the Client Access Point dialog box shown in Figure 15. The Client Access Point dialog box lets you name the CIFS role. I christened the Scale-Out CIFS with the name CIFS-Apps (see Figure 15). This is the server name that client applications use when they access the share. Clicking Next displays the Confirmation screen, which lets you confirm your selections or go back through the High Availability Wizard dialog boxes and make changes. If everything is OK, click Next on the Confirmation screen to display the Configure High Availability dialog box, which shows the progress of the CIFS configuration process. When it's complete, a Summary screen is displayed. Clicking Finish on the Summary screen closes the High Availability Wizard and returns you to the Failover Cluster Manager.

Figure 15
Client Access Point for
Scale-Out File Server



The next step is to add a file share to the CIFS scale-out application server. To create a new file share for the CIFS role, select the *Add File Share* link from the Actions pane, as I did for the general purpose file share in Figure 7. Clicking the *Add File Share* link for the scale-out CIFS starts the New Share Wizard shown in Figure 16.

To create a scale-out CIFS on the Select Profile dialog box, highlight the *SMB Share—Applications* option from the File share profile list

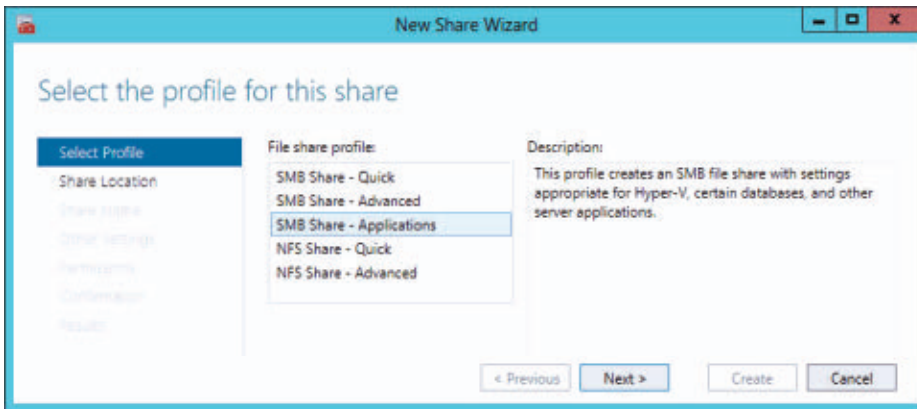


Figure 16
Selecting Profile for
Scale-Out File Server

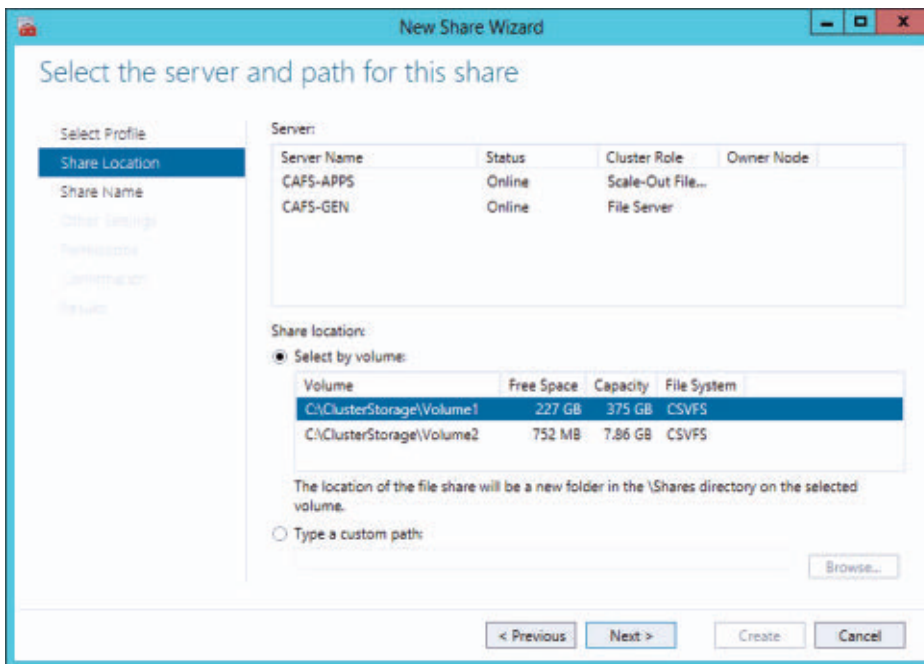


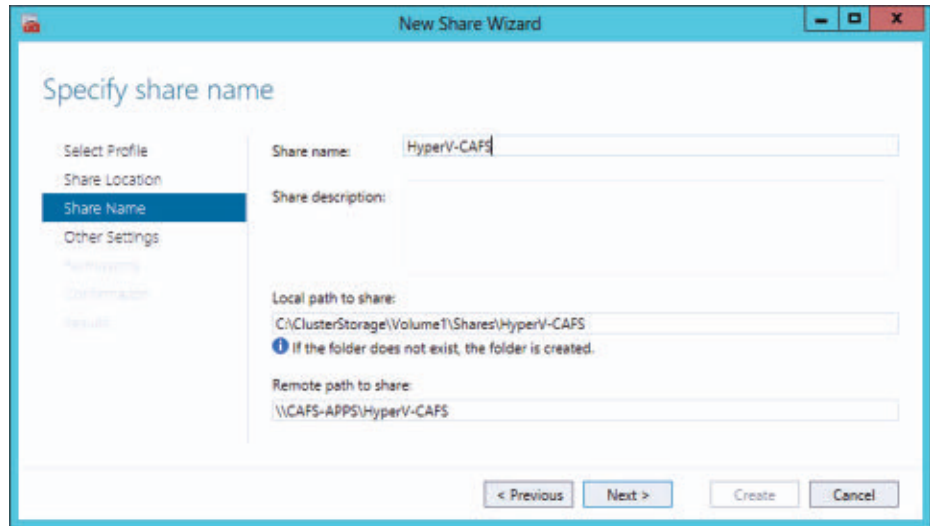
Figure 17
Share Location for
Scale-Out File Server

and then click Next to display the Share Location dialog box shown in Figure 17. The *Server* box near the top of the dialog box lists two CAFS file servers that were previously created. To add the CAFS to the scale-out application file server, select the CAFS-APPS file server that shows Scale-Out File Server in the Cluster Role column. Then select the CSV on which you want the CAFS share created.

This example has two existing CSVs. I selected C:\ClusterStorage\Volume1 as the location for the new scale-out CAFS. You also can enter a custom path to another CSV. After selecting the CSV, click Next to display the Share Name screen shown in Figure 18.

Figure 18

Share Name for Scale-Out File Server



The Share Name dialog box enables you to provide a name for the file share. I used the name HyperV-CAFS for the scale-out application CAFS (see Figure 18). In the center of the screen you also can see the local and remote paths to the CAFS. The local path for this example is C:\ClusterStorage\Volume1\Shares\HyperV-CAFS. The share will be accessed by networked systems using the path \\cafs-apps\HyperV-CAFS. Clicking Next displays the *Configure share settings* dialog box shown in Figure 19.

When you create a scale-out CAFS, the *Enable continuous availability* setting is checked by default. In addition, the *Enable access-based enumeration* and *Allow caching of share* settings are disabled. You cannot select them. The only other optional setting that you can choose is the *Encrypt data access* setting. I kept the default settings (see Figure 19). Clicking Next displays the *Specify permissions to control access* dialog box shown in Figure 20.

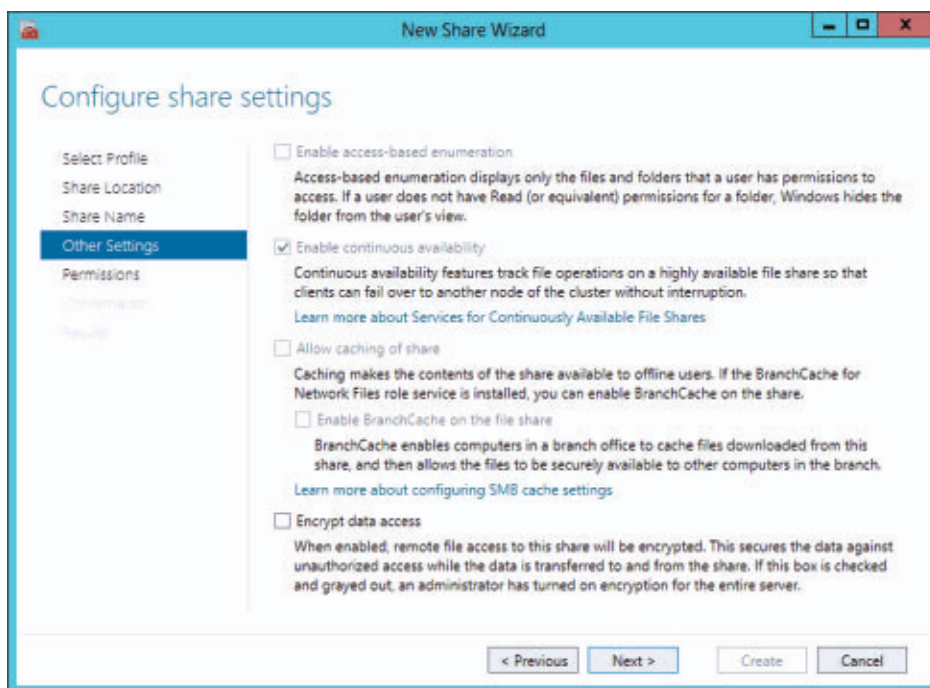


Figure 19
Configuring Share
Settings for the Scale-
Out File Server

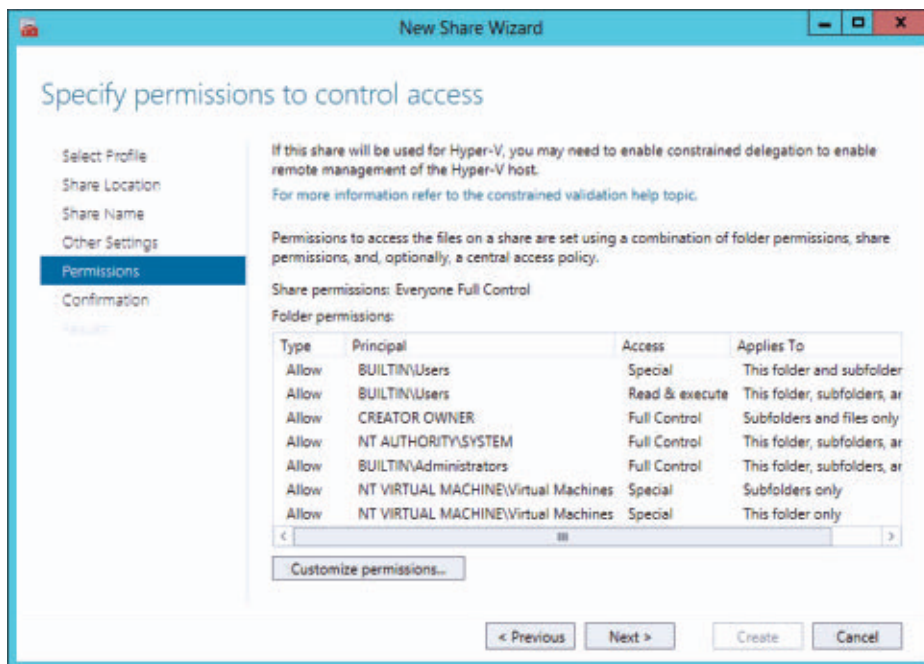
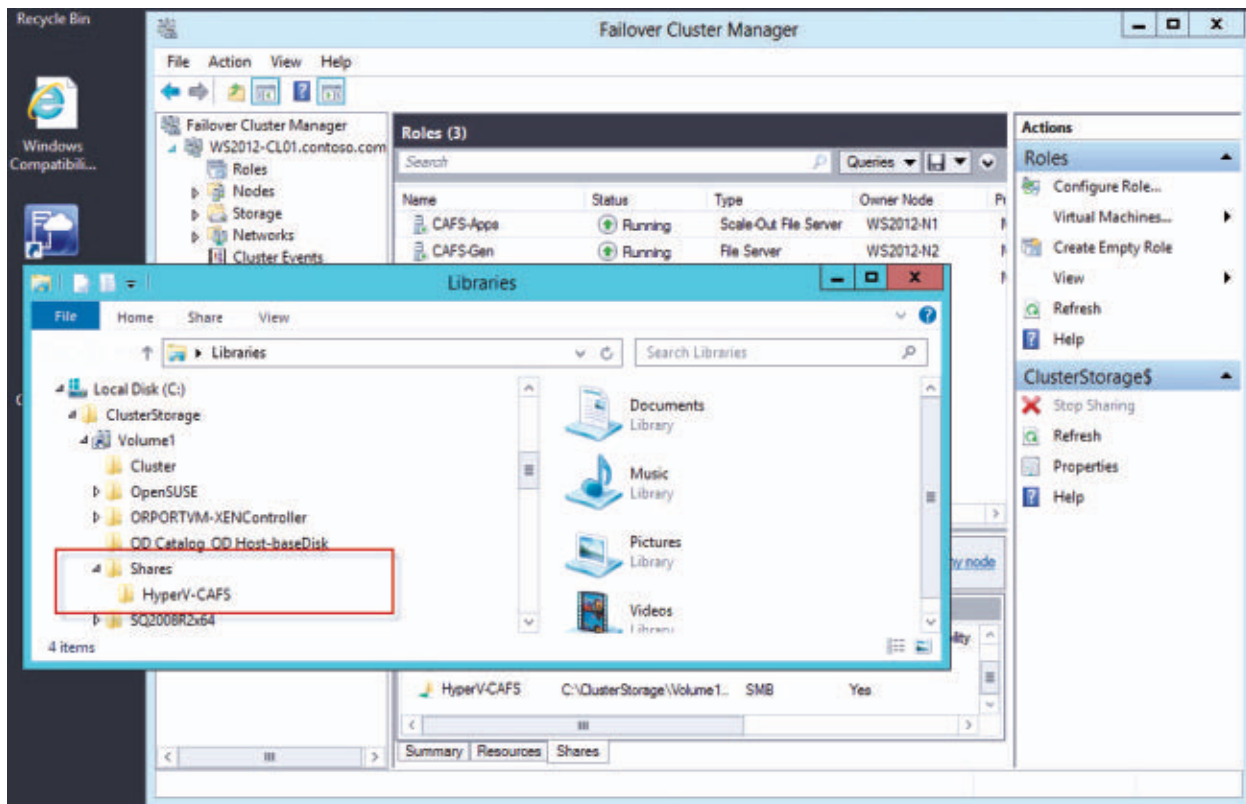


Figure 20
Specifying Permissions
for the Scale-Out File
Server

Like the general purpose CAFS, the scale-out CAFS is created with Full Control given to the Everyone group, which you'll probably want to change. I accepted the default permissions and clicked Next, which displays the Confirmation dialog box where you can see a summary of the choices that you made in the previous New Share Wizard dialog boxes. You can click Previous to go back and change any of the settings. Clicking Create on the Confirmation dialog box creates the scale-out CAFS and sets its permissions. After the share is created, it can be accessed locally from C:\ClusterStorage\Volume1\Shares\HyperV-CAFS or remotely from \\cafs-apps\HyperV-CAFS. The new CAFS is visible in the CSV mount point (Figure 21). At this point you can populate the share with Hyper-V VMs, SQL Server data, and log files or other types of application data.

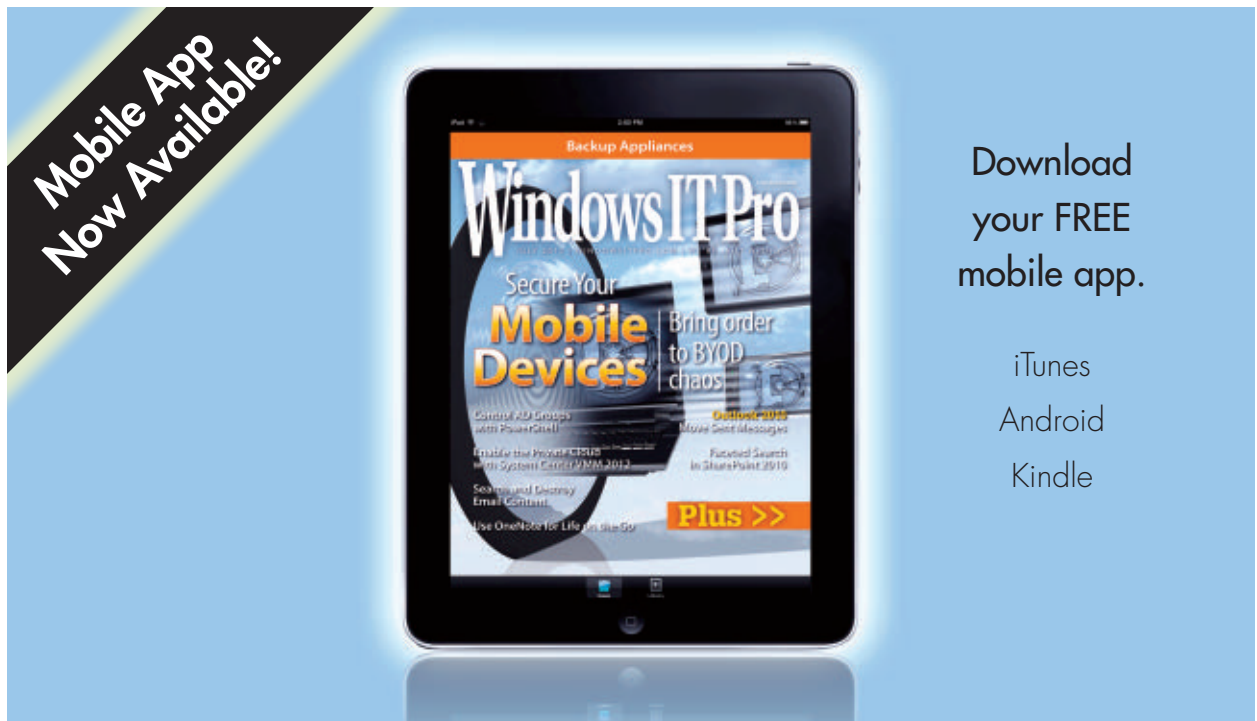
Figure 21

Accessing the CAFS
Share Locally



Improve File Availability

In this article I demonstrate how you can make use of CAFS to add increased availability and flexibility to your IT infrastructure. CAFS provides improved availability for general purpose file shares and also enables server applications such as SQL Server and Hyper-V to store their data on highly available file shares, increasing the range of storage options for your mission-critical applications. ■



**Mobile App
Now Available!**

Backup Appliances

Windows IT Pro

Secure Your Mobile Devices

Bring order to BYOD chaos

October 2013

Plus >>

Download your **FREE** mobile app.

iTunes
Android
Kindle

FAQ

Answers to Your Questions



John Savill



Jan De Clercq

Q: When should I use Resilient File System with Windows Server 2012 and Windows Server 2012 R2?

A: Windows Server 2012 introduced Resilient File System (ReFS) as an additional file system option. It features improved resiliency and availability over NTFS. However, in Windows Server 2012 and Windows Server 2012 R2 it also lacks some capabilities, which means it's not an option for many workloads (including SQL Server, Hyper-V, and many file server roles). The question is often asked: What, then, should I use ReFS for?

In Server 2012 and Server 2012 R2, it's actually fairly simple: Use ReFS for archived data. If you have critical data that needs to be archived and needs the highest levels of resiliency, such as for huge image files, archived VHD files, or anything else important, then store it on ReFS.

—John Savill

Q: Is there a Microsoft website that displays statistics about Microsoft solutions?

A: I stumbled across the [Microsoft by the Numbers](#) site, which is very interesting. It's a web page styled like the Windows Start screen, showing statistics about many of the major solutions from Microsoft, including Windows 8, SkyDrive, Windows Phone, Yammer, and Exchange. Take a look!

—John Savill

Q: What are some common best practices for securing the default Administrator account in a Windows Active Directory domain?

A: A common security best practice for protecting the Administrator account is to disable it, rename it, then change the text in its Description field. Not only does this hide the account, but it also hides the most visible indications that this is the almighty Administrator account. (However, you can recognize the Administrator account from its security identifier—SID, which ends in 500.)

Another option is to create a decoy user account called Administrator that has a very limited set of permissions or no special permissions or user rights. If you don't want to disable the Administrator account (it can be a life-saver if you lock out your day-to-day admin account), it's a good idea to always give the account a long, complex, and random password that you change at regular intervals.

Finally, make sure that you have an automated procedure in place to accomplish these tasks. For automation, use a combination of Group Policy Object (GPO) settings and Windows PowerShell scripts.

—Jan De Clercq

Q: Why are the variables in my Windows PowerShell scripts not working as I expected in other strings?

A: When you create a variable such as \$var, the way PowerShell knows it's a variable is that it starts with the dollar sign. In some circumstances, PowerShell can automatically translate the variable into its actual value as part of another string. For example, I could type the following and get the output you see on the third line:

```
$var = "John"
write-host "Hello $var"
Hello John
```

But if I try to use a single quote around *Hello \$var*, it doesn't work:

```
$var = "John"
write-host 'Hello $var'
Hello $var
```

Double-quoted (or double quotation mark) strings expand environment variables, and single-quoted (or single quotation mark) strings do not. An alternative is to use PowerShell commands such as `-join`. Or you could concatenate strings by using `+`. You could even use the following:

```
$var = "John"
$var2 = "Savill"
"Hello {0} {1}" -f $var,$var2
```

The problem is even bigger if the value of your variable is an object that has its own attributes. For example, take a look at this:

```
$notepadproc = get-process notepad
write-host "The process ID of Notepad is $notepadproc.Id"
The process ID of Notepad is System.Diagnostics.Process
(notepad).Id
```

This is clearly not what I wanted. The problem is, only the variable gets expanded in a string, not property extensions. This is why anything after the variable name is output as part of the string. The solution is to put the whole expression into brackets (aka parentheses):

```
$notepadproc = get-process notepad
write-host "The process ID of Notepad is $($notepadproc.Id)"
The process ID of Notepad is 4640
```

—John Savill

Q: What exactly are the Virtual Smart Cards that Microsoft supports in Windows 8? How are they different from traditional physical smart cards?

A: Virtual Smart Cards (VSCs) let users with a Windows 8 computer equipped with a Trusted Platform Module (TPM) chip that meets the TPM 1.2 specification leverage the benefits of physical smart card logon without making an investment in smart card hardware and without the possibility of losing a card. Windows 8 VSCs are based on a software construct that emulates a smart card on the OS level. VSCs appear to Windows 8 the same way they would as physical smart cards, and they use the same application-level APIs. For a user, logging on with a VSC is as easy as logging on with a password; all he has to do is enter his PIN (there's no need to insert a physical card in a card reader or connect a USB token).

Like traditional physical smart cards, VSCs provide a two-factor authentication mechanism. Physical smart cards are physical objects and clearly provide a “something you have” authentication factor. With VSCs there is also always a hardware element involved: the TPM. Just like physical smart cards, VSCs are always used in conjunction with a “something you know” (e.g., a password or a PIN) authentication factor to complete the two-factor authentication.

VSCs are secure because even though the private keys the VSC holds are physically stored on the computer's hard drive, the keys are encrypted using a secret that is securely stored on the TPM, which is tamperproof. A direct consequence of using the TPM is that you can't move a VSC to a different computer. This is because only a local machine's TPM that encrypted the keys is able to use them. That also means users can't use the same VSC from multiple machines and attackers can't remove the hard drive to get access to the VSC and its private keys. This non-exportability is also an important security characteristic of physical smart cards: The information stored on a physical card can't be extracted to be used somewhere else.

Windows 8 and its applications see a VSC as being always inserted in a virtual card reader. This means that unlike with physical smart cards, administrators can't set a policy to automatically log the user off when the card is removed. Like physical smart cards, VSCs will lock out a user who enters an incorrect PIN a specified number of times.

You can find more information on Windows 8 VSCs in the “[Understanding and Evaluating Virtual Smart Cards](#)” white paper.

—Jan De Clercq

Q: How can I see every Windows Azure image available?

A: Although many available images are shown in the Windows Azure IaaS VM creation wizards, you can actually view many more. To list all of them, use Windows PowerShell. After you configure your machine with the Windows Azure cmdlets and configure your connection, enter the PowerShell command below:

```
Get-AzureVMImage | ft Label,ImageName,LogicalSizeInGB
```

—John Savill

Product News for IT Pros

Veeam Backup & Replication 7.0 Debuts

Veeam Software released Veeam Backup & Replication 7.0. The solution introduces two exciting innovations—Built-in WAN Acceleration and Backup from Storage Snapshots—that take Veeam’s Modern Data Protection to the next level. Developed by Veeam and optimized specifically for Veeam backups, Built-in WAN Acceleration copies data to offsite locations up to 50 times faster than a regular file copy. It also eliminates the need to purchase and deploy a general-purpose WAN acceleration appliance or acquire additional network bandwidth for offsite backups. Developed in partnership with HP, Backup from Storage Snapshots dramatically improves recovery point objectives (RPOs) and greatly reduces stress on the virtual infrastructure. As a result, IT admins can make backups as often as they want, even for I/O-intensive virtual machines (VMs). To learn more, visit the [Veeam Software website](#).



Nimble Storage Leverages Cisco UCS and VMware vSphere

Nimble Storage announced a new converged infrastructure reference architecture leveraging Cisco Unified Computing System (UCS) and VMware vSphere. At the core of this converged infrastructure are VMware vSphere 5.1, flash-optimized Nimble Storage CS-Series arrays, and Cisco UCS B-Series blade servers. Together, these components allow IT organizations to scale as needed to respond faster to changing business needs, while lowering project risks and capital expenses. Nimble Storage launched a series of SmartStack pre-validated reference architectures that minimize the challenges associated with deploying application, server, hypervisor, networking, and storage components as an integrated solution. At the heart



of every SmartStack solution is a pre-validated reference architecture leveraging Nimble Storage CS-Series and strategic partners' solutions. By testing and validating the combined solutions, Nimble Storage creates prescriptive architectures that help organizations accelerate deployment and minimize risk associated with deploying solutions in the data center, whether it be for server virtualization, virtual desktop infrastructure (VDI), or data protection. For more information, check out the [Nimble Storage website](#).



NCP Enhances Windows VPN Client and Gateway

NCP announced that it has made new versions of its Windows-compatible IPsec VPN client suite and hybrid IPsec/SSL VPN gateway available to the channel. The key features in version 9.32 of the NCP Secure Enterprise Client and version 8.11 of the NCP Secure Enterprise VPN Server were designed to help enterprise customers yield a higher level of security while ensuring maximum performance during remote access sessions. NCP's VPN client suite and gateway now come equipped with support for elliptic curve cryptography (ECC) to safeguard VPN connections. Public-key cryptography based on ECC currently offers both higher security and better performance compared with RSA. To further boost remote access performance, the NCP Secure Enterprise VPN Server offers optimized multi-processor support. For more information, visit the [NCP website](#).



GlobalSign Provides Automated Certificate Lifecycle Management

GlobalSign announced the availability of the GlobalSign Auto Enrollment Gateway. AEG integrates with Active Directory (AD), allowing enterprises to automate the enrollment, provisioning, and management of GlobalSign digital certificates for Windows environments. By replacing their internal CAs with GlobalSign's services, enterprises strengthen security and reduce costs by adding certificate-based solutions such as two-factor authentication and advanced SSL without

having to manage their own highly complex and costly internal certificate authority (CA). Because GlobalSign SaaS certificate services provide the latest best practices and highest standards in certificate technology, enterprises using them reduce their risk of falling victim to attacks that take advantage of weak and mismanaged certificates. Eliminating the need to manage a resource-intensive internal CA reduces the total cost of ownership (TCO) of the public key infrastructure (PKI) as well as the risk of system outages that stall business activities. For more information, visit the [GlobalSign website](#).

NETIKUS.NET Expands EventSentry Light

NETIKUS.NET announced a major update to EventSentry Light, the free edition of its real-time monitoring solution, EventSentry. The product now offers significantly more functionality that was previously available only in the commercial edition of EventSentry. EventSentry Light can monitor up to two Windows-based computers and two network devices. Users who need to monitor more hosts or require more functionality can seamlessly upgrade to the full version of EventSentry. The full edition of EventSentry includes web-based reporting, log consolidation, compliance tracking functionality, and more, all backed by the company's acclaimed customer support. Support for EventSentry Light is available through the NETIKUS.NET forums, which are continuously monitored by NETIKUS.NET staff. You can download EventSentry Light from the [NETIKUS.NET website](#).



CommVault Stops Virtual Machine Sprawl

CommVault announced the industry's first virtual machine (VM) intelligent archiving capability to help enterprises and service providers eliminate VM sprawl and regain control of virtual infrastructure resources. VM sprawl results from pervasive deployment and growth of VMs, some of which then sit unutilized long after their useful lives. CommVault Simpana VM Archiving reclaims idle virtual host



and shared storage resources by automatically managing and moving unused VMs to cost-effective storage, with the ability to instantly recover archived VMs, for increased utilization, efficiency, and savings. Available at no additional cost to customers who use Simpana 10 software for virtual server protection, VM Archiving is integrated with CommVault's singular software platform, which enables users to instantly deploy, protect, and archive virtual and physical servers from a single management console. The new intelligent archive capability is a key component of CommVault's modern data-management approach to continuously address operational and protection concerns in virtualized and Infrastructure as a Service (IaaS) environments. For more information about VM Archiving, visit the [CommVault website](#).



ENow Management System 6.0 Delivers Exchange Server 2013 and Lync Support

ENow announced the release of EMS 6.0, which promises to make the jobs of Exchange Server and Lync administrators easier while increasing service availability. The EMS 6.0 release enables Exchange 2013 administrators to proactively monitor and achieve visibility into their messaging infrastructure. EMS 6.0's Mailscape module proactively tests all the core messaging components, including DAG configuration, external and internal mail flow, OWA, and ActiveSync. The reporting module has more than 210 reports, including detailed insight on mobile device usage. Also included is a new module, named UniScope, that provides visibility into Microsoft Lync deployments. UniScope proactively tests the core components of a Lync deployment, including web conferencing, mediation servers, end user connectivity, and address book downloads. For more information, visit the [ENow website](#). ■

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowssitpro.com

Support
Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowssitpro.com

News
Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS
Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS
Windows IT Pro VIP
Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowssitpro.com/vip-premium-membership

SQL Server Pro
Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro
Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro
Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

Windows IT Pro..... 1, 2, 39, 69

Vendor Directory

CommVault..... 77, 78

ENow..... 78

Facebook..... 19, 20

GlobalSign..... 76, 77

Google..... 20

LinkedIn..... 19, 20

NCP..... 76

NETIKUS.NET..... 77

Nimble Storage..... 75, 76

Veeam Software..... 75

